LA-UR-16-20786

Approved for public release; distribution is unlimited.

Title:	Expanding Multichannel Waveform Correlation Detectors to Target Template-Dissimilar Waveforms (Part I)
Author(s):	Carmichael, Joshua Daniel
Intended for:	Bulletin of the Seismological Society of America
Issued:	2016-02-29 (rev.1)

Expanding Multichannel Waveform Correlation Detectors to Target Template-Dissimilar Waveforms (Part I)

Joshua D Carmichael

Los Alamos National Laboratory

Bikini Atoll Rd., SM 30

Los Alamos, NM 87544, USA

February 12, 2016

Abstract

Waveform correlation detectors used in seismic monitoring scan geophysical data to test two competing hypotheses: that data contain (1) a noisy, amplitude-scaled version of a template waveform, or, (2) only noise. In reality, geophysical wavefields include signals triggered by non-target sources (background seismicity) and target signals that are only partially correlated with the waveform template. We reform the waveform correlation detector hypothesis test to accommodate deterministic uncertainty in template/target waveform similarity and thereby derive a new detector using convex set projections for use in explosion monitoring. Our analyses give probability density functions that quantify the detectors' degraded performance with decreasing waveform similarity. We then apply our results to three announced North Korean nuclear tests and use International Monitoring System (IMS) arrays to determine the probability that low magnitude, off-site explosions can be reliably detected with a given waveform template.

Introduction

Seismic sources with colocated hypocenters that are well separated in time often produce correlated waveforms. Monitoring missions exploit this waveform similarity to identify repeated explosions from localized target regions with multichannel correlation detectors (MacCarthy et al. (2008); Gibbons and Ringdal (2006); Carmichael et al. (2016)). These detectors operate as matched filters by scanning reference, or “template” waveforms, recorded by multi-element receiver networks against commensurate data streams to search for signals of identical shape with unknown amplitude. Significant correlation between such a network-measured template and a potential waveform match (target waveform) thereby requires similar source time functions and locations. This means that a matching template and target waveform include equivalent features within their wavefronts that record identical relative arrival times at each array receiver (same raypaths). Correlation detectors quantify evidence of such matching by comparing a wavefield-derived detection statistic to a noise-dependent declaration threshold (Harris (1989, 1991)). The significance of the detection statistic, its threshold, and their relation to template/target waveform similarity is conditional on several factors that include template bandwidth and target signal amplitude (Weichecki-Vergara et al. (2001); Ford and Walter (2015)). Deterministic uncertainties in these signal attributes, however, often ambiguously relate to source attributes such as spatial separation or relative magnitude of target explosions (Dodge and Walter (2015)). For instance, signals triggered by a large explosion that is spatially separated from a template source can produce the same correlation value as waveforms triggered by a relatively weak, template-colocated explosion (Schaff (2010)). Because these correlation statistics are indistinguishable, certain estimation methods will erroneously produce identical relative magnitude estimates for both event-pairs. In these cases, uncertainty in correlation parameters can degrade the reliability of related source parameter estimates.

These correlation challenges have practical consequences for monitoring agencies like the Comprehensive Nuclear Test Ban Treaty Organization (CTBTO), who are charged with detecting signatures of nuclear explosions. To perform this mission, the CTBTO assembles and processes data collected from a global network of instruments called the International Monitoring System (IMS) that includes seismic arrays. Waveform processing pipelines at the CTBTO, in particular, implement correlation detectors using these array data to identify repeating explosions at known nuclear test sites (Bobrov et al. (2014)). It is therefore crucial to the CTBTO mission that they understand the operational performance of these detectors in realistic monitoring scenarios, such as when sources are separated by distances that are comparable to a seismic wavelength. The importance of establishing such performance limits was recently highlighted by alleged seismic evidence that supported claims of an unannounced nuclear test in North Korea during May of 2010 (Zhang and Wen (2015)). In that case, researchers declared a correlation match between a low magnitude ($m_b < 2$) seismic event and template waveforms measured from prior North Korean tests, at a statistical threshold conventionally considered uninterpretable. Their hypocentral solutions also placed the supposed explosion source away from previous test locations. These claims motivated seismologists to reevaluate the general applicability of using waveform correlation with IMS data to identify low magnitude explosions near test sites. Ford and Walter ((2015)) adapted a correlation detector to spectrally rescale target explosions by a specific source size and thereby monitor the North Korean Test Site (NKTS) for waveforms triggered by explosions with higher-than-template corner frequencies. Carmichael and Hartse (Carmichael and Hartse (2016)) modified a similar correlation detector to account for background seismicity and template/target source separation, and thereby estimated threshold magnitudes for explosions located near the NKTS. While these studies provide insight into correlation, they also raise additional questions. In particular, it remains unclear if different detectors could model unanticipated inconsistencies between a template and a hypothetical target waveform's geometry. More research is therefore required to understand the influence of waveform dissimilarity on the capability of waveform detectors that process IMS data. In this paper, we quantify how deterministic uncertainties in target waveform geometry superimpose with noise to inflate the magnitude at which a waveform is reliably detected with IMS arrays. In particular, we develop a new detector using convex set projections to compare template and target signal decorrelation against the probability of explosion detection. We then apply this detector to three announced North Korean nuclear tests recorded on IMS stations and empirically estimate detection probabilities for a range of source sizes and hypocentral separation distances. We thereby construct receiver operating characteristic (ROC) curves to demonstrate that our convex-set detector performs competitively with a correlation detector while showing a better agreement with predictions and producing fewer empirical false alarms. We suggest that this detector provides a more realistic assessment of IMS monitoring capability when target explosions are not collocated with template sources, or waveforms are not entirely repeatable. While our analyses focus on explosion monitoring, our methods are applicable to geophysical waveform detection from any source type.

Correlation Detectors

We write ground motion data $x_k(t)$ that is recorded on channel k of a seismic network and digitized from time t_0 at interval Δt until time $T = N \cdot \Delta t$ as:

$$\boldsymbol{x}_k = \Big[x_k(t_0)\,,\, x_k(t_0 + \Delta t)\,,\, \cdots\,,\, x_k(t_0 + (N-1)\Delta t) \Big]^{\mathrm{T}}. \tag{1}$$

We similarly represent N-sample, L-channel data as matrix $\boldsymbol{x} = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_L]$, where column k contains signal $\boldsymbol{x}_k$ sampled from time t_0 until time $t_0 + (N-1)\Delta t$, and row l of $\boldsymbol{x}$ contains a network-wide sample of ground motion at $t_0 + (l-1)\Delta t$. We use analogous notation for multichannel template waveforms $\boldsymbol{w}$ and noise $\boldsymbol{n}$. To find signals within a noisy data stream with the same shape as such a template, we compare two competing hypotheses. The first hypothesis $\mathcal{H}_0$ presumes that data $\boldsymbol{x}$ contain only zero mean Gaussian noise ($\boldsymbol{n}$) of unknown variance. The second hypothesis ($\mathcal{H}_1$) presumes $\boldsymbol{x}$ consists of an amplitude scaled template waveform ($\boldsymbol{w}$) of unknown amplitude A plus Gaussian noise of unknown variance ($\boldsymbol{x} = A\boldsymbol{w} + \boldsymbol{n}$):

$$\begin{aligned} &\mathcal{H}_0 : \boldsymbol{x} = \boldsymbol{n} \sim \mathcal{N}\left(\boldsymbol{0}, \sigma^2 \boldsymbol{I}\right) \\ &\qquad \text{(noise present, } \sigma \text{ unknown)} \\ &\mathcal{H}_1 : \boldsymbol{x} = A\boldsymbol{w} + \boldsymbol{n} \sim \mathcal{N}\left(A\boldsymbol{w}, \sigma^2 \boldsymbol{I}\right), \\ &\qquad \text{(noise present, } A, \sigma \text{ unknown)} \end{aligned} \tag{2}$$

where $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ symbolizes the multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. Temporal sample correlation within a channel or between sensor channels generally requires a non-diagonal covariance structure so that $\boldsymbol{\Sigma} \neq \sigma^2 \boldsymbol{I}$, in contrast to Equation 2. However, such structure is often representable using a scalar N_E that parametrizes each probability density function (PDF) and represents the effective number of statistically independent samples in $\boldsymbol{x}$; Appendix B outlines estimating N_E as $\hat{N}_E$ in the context of correlation.

We derive a multichannel correlation detector from the competing hypotheses in Equation 2 with a generalized likelihood ratio test for $\boldsymbol{x}$. In simplest terms, this ratio divides the PDF of $\boldsymbol{x}$ under $\mathcal{H}_1$ by the PDF for $\boldsymbol{x}$ under $\mathcal{H}_0$, where each function is evaluated at the maximum likelihood estimates (MLEs) of its unknown parameters (σ and A). The resultant statistic $r(\boldsymbol{x})$ gives an estimate for the true waveform cross correlation ρ. The associated ratio test on $r(\boldsymbol{x})$ compares the similarity between a template $\boldsymbol{w}$ and commensurate data stream $\boldsymbol{x}$ against a noise-dependent threshold η according to the decision rule (Harris (1991)):

$$r(\boldsymbol{x}) = \frac{\sum_k^L \boldsymbol{x}_k^{\mathrm{T}} \boldsymbol{w}_k}{\sqrt{\sum_k^L \boldsymbol{x}_k^{\mathrm{T}} \boldsymbol{x}_k \cdot \sum_k^L \boldsymbol{w}_k^{\mathrm{T}} \boldsymbol{w}_k}} \triangleq \frac{\langle \boldsymbol{x}, \boldsymbol{w} \rangle}{\|\boldsymbol{x}\| \|\boldsymbol{w}\|} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta, \tag{3}$$

where the Frobenius inner product $\langle \boldsymbol{x}, \bullet \rangle$ and Frobenius norm $\| \bullet \|$ extend vectorial dot product operations to matrix data. The Neyman-Pearson criteria determines a particular threshold η for event detection from the PDF of $r(\boldsymbol{x})$. This criteria uses a constant false-alarm on noise constraint $\Pr_{FA}$ to compute η as the inverse, right-tail probability of $r(\boldsymbol{x})$ under the null hypothesis $\mathcal{H}_0$:

$$\Pr_{FA} = \int_\eta^1 f_R(r; \mathcal{H}_0)\, dr, \tag{4}$$

where $f_R(r; \mathcal{H}_0)$ is the PDF for $r(\boldsymbol{x})$ under $\mathcal{H}_0$. This threshold also sets the target waveform detection probability Pr_D:

$$\mathrm{Pr}_D = \int_{\eta}^{1} f_R\left(r; \mathcal{H}_1\right) dr \tag{5}$$

where $f_R\left(r; \mathcal{H}_1\right)$ is the PDF for $r\left(\boldsymbol{x}\right)$ under $\mathcal{H}_1$. We derive the general PDF under hypothesis $\mathcal{H}_k$ $(k = 0, 1)$ in Appendix A (Equation A.6), where we show:

$$\begin{aligned} f_R\left(r; \mathcal{H}_k\right) = \mathrm{B}&\left(r^2; \frac{1}{2}, \frac{1}{2}\left(N_E - 1\right), \lambda, \lambda^{\perp}\right) + \\ \mathrm{B}&\left(-r^2; \frac{1}{2}, \frac{1}{2}\left(N_E - 1\right), \lambda, \lambda^{\perp}\right). \end{aligned} \tag{6}$$

Here $\mathrm{B}\left(t, N_1, N_2, \alpha, \beta\right)$ is the doubly noncentral Beta distribution function. It is evaluated at t, has N_1 and N_2 degrees of freedom, and noncentrality parameters α and β. Given these parameterizations, Equation 3 with Equation 4 define the multichannel correlation detector and Equation 5 quantifies its detection performance. The scalars λ and $\lambda^{\perp}$ shaping $f_R\left(r; \mathcal{H}_k\right)$ are respectively proportional to the template-coherent and template-incoherent portions of the target waveform energy (Equation A.3). Therefore, λ quantifies the detection power of the correlation detector (Figure 1). When $\lambda^{\perp} = \lambda = 0$, no signal is present and $\mathcal{H}_0$ is true. When $\lambda^{\perp} = 0$ and $\lambda > 0$, $\mathcal{H}_1$ is true. Then the target and template waveforms are identical at each network receiver except in amplitude and noise content. In such cases, correlation detectors provide an optimal capability to identify waveforms of known shape in Gaussian noise (Kay (1998, page 133)). In more realistic cases, the seismic wavefield includes more than just noise, or waveforms of known shape. Consequently, the hypotheses in Equation 2 misrepresent observations. The null hypothesis insufficiently describes noisy signals like $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ $(\boldsymbol{u} \neq \boldsymbol{0})$ that are triggered by background seismicity (earthquakes, mining blasts). Such waveforms may be partially correlated with the template and appear to satisfy $\mathcal{H}_1$. Likewise, the alternative model/hypothesis in Equation 2 insufficiently describes noisy target waveforms $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ $(\boldsymbol{u} \neq A\boldsymbol{w})$ that are partially decorrelated with the template. Such target events may be triggered by explosions of interest that are spatially separated from the template's source, so that phase segments within each wavefront take dissimilar paths and destructively interfere. These signals would be partially incoherent with the template and may appear to satisfy $\mathcal{H}_0$. We write this latter hypothesis error symbolically:

$$\begin{aligned} \mathcal{H}_1 \text{ error: } \boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n} \neq A\boldsymbol{w} + \boldsymbol{n} \\ \implies \text{template/target mismatch,} \end{aligned} \tag{7}$$

and consider $\mathcal{H}_0$ errors triggered by template-correlated background seismicity in parallel work. Proper treatment of realistic errors under $\mathcal{H}_1$ requires waveform detection methods that account for scenarios where target signals are not limited to amplitude-scaled copies of a template.

Figure 2 illustrates two competing cases where single-channel target waveforms show different levels of similarity with a template. The left panel shows an example in which the target signal is reasonably represented with the conventional correlation hypothesis (Equation 2). Specifically, peaks and troughs within the template and target waveforms generally match in timing and amplitude (Figure 2a) so that localized cross-correlation is reasonably stable over the high energy portions of the signal (Figure 2b). The right panel provides an example where the correlation hypothesis poorly represents amplitude scaling between the template and target waveform (Equation 7). In this latter case, the relationship between peaks and troughs of the template and target waveform (Figure 2c) is often unclear. The localized cross-correlation alternatively shows that some waveform segments have moderate correlation, while others are entirely decorrelated (Figure 2d).

Estimating Deterministic Decorrelation

To quantify decorrelation caused by waveform dissimilarity, we consider the correlation coefficient between a zero-mean template waveform $\boldsymbol{w}$ and a zero-mean target waveform $\boldsymbol{u}$. Both waveforms are respectively contaminated by zero-mean noise $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ of variance σ_1^2 and σ_2^2 that is stationary over their N-sample lengths. The true correlation coefficient ρ between these waveform data is defined as:

$$\rho = \frac{\mathbb{E}\left\{\langle \boldsymbol{u} + \boldsymbol{n}_1, \boldsymbol{w} + \boldsymbol{n}_2 \rangle\right\}}{\sqrt{\mathbb{E}\left\{||\boldsymbol{u} + \boldsymbol{n}_1||^2\right\} \mathbb{E}\left\{||\boldsymbol{w} + \boldsymbol{n}_2||^2\right\}}}, \tag{8}$$

where $\mathbb{E}\left\{\bullet\right\}$ is the expected value operator. Equation 8 is reducible to a product of two distinct coefficients that represent different sources of degradation from perfect correlation. One is attributable to the effects of noise. The other is attributable to the dissimilarity between the underlying noise-free waveforms. We illustrate this factorization by first computing certain moments of the noise field to simplify the following algebra: since $\mathbb{E}\left\{||\boldsymbol{n}_k||^2\right\} = \sigma_k^2(N_E - 1)$, and $\mathbb{E}\left\{\langle \boldsymbol{w}, \boldsymbol{n}_k \rangle\right\} = \mathbb{E}\left\{\langle \boldsymbol{u}, \boldsymbol{n}_k \rangle\right\}$ $= 0$ ($k = 1,\, 2$), it follows from Equation 8 that:

$$\rho = \frac{||\boldsymbol{w}||\,||\boldsymbol{u}||}{\sigma_1\,\sigma_2} \cdot \frac{\dfrac{1}{(N_E - 1)^2}\dfrac{\mathbb{E}\left\{\langle \boldsymbol{w}, \boldsymbol{u} \rangle\right\}}{||\boldsymbol{w}||\,||\boldsymbol{u}||}}{\sqrt{\mathbb{E}\left\{\dfrac{||\boldsymbol{u}||^2}{||\boldsymbol{n}_1||^2} + 1\right\} \mathbb{E}\left\{\dfrac{||\boldsymbol{w}||^2}{||\boldsymbol{n}_2||^2} + 1\right\}}}. \tag{9}$$

We then factor ρ into a product of three terms:

$$\rho = \frac{\dfrac{||\boldsymbol{u}||}{\sigma_1(N_E - 1)}}{\sqrt{\mathbb{E}\left\{\dfrac{||\boldsymbol{u}||^2}{||\boldsymbol{n}_1||^2} + 1\right\}}} \frac{\dfrac{||\boldsymbol{w}||}{\sigma_2(N_E - 1)}}{\sqrt{\mathbb{E}\left\{\dfrac{||\boldsymbol{w}||^2}{||\boldsymbol{n}_2||^2} + 1\right\}}} \frac{\langle \boldsymbol{w}, \boldsymbol{u} \rangle}{||\boldsymbol{w}||\,||\boldsymbol{u}||} \tag{10}$$

and rewrite these factors using the definition of SNR:

$$\begin{aligned} \rho &= \underbrace{\frac{\sqrt{\mathrm{SNR}\,(\boldsymbol{w})}}{\sqrt{\mathrm{SNR}\,(\boldsymbol{w}) + 1}} \frac{\sqrt{\mathrm{SNR}\,(\boldsymbol{u})}}{\sqrt{\mathrm{SNR}\,(\boldsymbol{u}) + 1}}}_{\text{noise part: } \rho_0} \underbrace{\frac{\langle \boldsymbol{w}, \boldsymbol{u} \rangle}{||\boldsymbol{w}||\,||\boldsymbol{u}||}}_{\text{signal part: } \rho_\infty} \\ &= \rho_0 \cdot \rho_\infty \end{aligned} \tag{11}$$

where SNR($\boldsymbol{w}$) abbreviates the ratio of signal to noise power of $\boldsymbol{w}$. The "noise" correlation term ρ_0 in Equation 11 depends on waveform SNR, whereas ρ_∞ measures the deterministic correlation between the underlying infinite SNR signals; we equivalently call $1-\rho_\infty$ the deterministic waveform dissimilarity. In the case that $\boldsymbol{u} = A\boldsymbol{w}$, $\rho_\infty = 1$, and the true correlation equates the noise correlation ($\rho = \rho_0$). In all other cases, deterministic differences between the template and the detected waveforms contradict $\mathcal{H}_1$, imply $\rho_\infty < 1$, and indicate the repeating-signal hypothesis is not representative of the detected waveform. To estimate ρ_∞ using the sample correlation $r(\boldsymbol{x})$, we rearrange Equation 11 and suggest the estimator $\hat{\rho}_\infty$:

$$\hat{\rho}_\infty = \text{sign}\left[r(\boldsymbol{x})\right] \cdot \min\left\{\frac{|r(\boldsymbol{x})|}{\rho_0},\ 1\right\} \tag{12}$$

where $\min\{\bullet\}$ prevents $\hat{\rho}_\infty$ from assuming nonsensical correlation values that exceed one. Scalar ρ_0 is estimable from reference waveforms, as detailed in our error analysis (Appendix E).

Correlation Under an $\mathcal{H}_1$ Error

We now form a test on target waveforms that includes deterministic uncertainties in waveform correlation in the presence of background noise. Empirical detection routines typically show that the Gaussian model represents observed noise ($\boldsymbol{n}$) well (Carmichael and Hartse (2016); Carmichael et al. (2016)). The deterministic model ($A\boldsymbol{w}$) for the such signals, however, can poorly represent target waveforms. An admissible alternative hypothesis $\mathcal{H}_1$ must instead permit some deterministic incoherence between template and target waveforms without increasing false triggers on noise. Quantitatively, we expect non-unit correlation between a template $\boldsymbol{w}$ and noise-free portions of a target waveform. If such a target waveform is $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ (Equation 7), this means that:

$$\frac{\langle \boldsymbol{u}, \boldsymbol{w} \rangle}{||\boldsymbol{u}||\,||\boldsymbol{w}||} \triangleq \rho_\infty \leq 1, \tag{13}$$

where the ∞ subscript indicates the noise-free, or infinite SNR waveform correlation (Equation 11). Potential target waveforms $\boldsymbol{u}$ are then constrained to a set $\mathcal{C}$ that includes all signals that correlate above a signal-dependent threshold ρ_∞ with the template $\boldsymbol{w}$:

$$\mathcal{C} = \left\{ \boldsymbol{u} : \frac{\langle \boldsymbol{u}, \boldsymbol{w} \rangle}{||\boldsymbol{u}||\,||\boldsymbol{w}||} \geq \rho_\infty \right\}. \tag{14}$$

Equation 14 geometrically represents a high-dimensional cone in $N \times L$-dimensional space with a vertex parallel to $\boldsymbol{w}$ and aperture $\arccos(\rho_\infty)$ The set of all possible target waveforms thereby occupy the interior of $\mathcal{C}$ ($\mathcal{C}^\circ$) or its boundary ($\partial\mathcal{C}$). Set $\mathcal{C}$ then forms the signal present model of our reformed hypotheses test:

$$\begin{aligned} &\mathcal{H}_0 : \boldsymbol{x} = \boldsymbol{n} \sim \mathcal{N}\left(\boldsymbol{0}, \sigma^2 \boldsymbol{I}\right) \\ &\qquad \text{(noise present, } \sigma \text{ unknown)} \\ &\mathcal{H}_1 : \boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n} \sim \mathcal{N}\left(\boldsymbol{u}, \sigma^2 \boldsymbol{I}\right) \\ &\qquad \text{(noisy target present, } \boldsymbol{u} \in \mathcal{C},\ \sigma \text{ unknown).} \end{aligned} \tag{15}$$

We compare these hypotheses with an associated detection statistic that we derive from a generalized log-likelihood ratio on $\boldsymbol{x}$ (as with $r(\boldsymbol{x})$). This derivation (Appendix C) exploits the theory of projection-onto-convex sets (Stark and Yang (1998, pg. 111)) and shows that the "cone" statistic $s(\boldsymbol{x})$ depends on the size (norm) of the conic projection $P_{\mathcal{C}}(\boldsymbol{x})$ of $\boldsymbol{x}$ onto the boundary $\partial\mathcal{C}$ of $\mathcal{C}$. The resulting scalar test provides a decision rule similar to the test in Equation 3 that uses $r(\boldsymbol{x})$. The cone decision rule is:

$$s(\boldsymbol{x}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta, \tag{16}$$

with detection statistic:

$$s(\boldsymbol{x}) = \begin{cases} 0: & \frac{r}{\sqrt{1-r^2}} < -c, & P_{\mathcal{C}}(\boldsymbol{x}) = \boldsymbol{0} \\ \frac{\gamma}{||\boldsymbol{x}||}: & \frac{r}{\sqrt{1-r^2}} \in \left[-c, \frac{1}{c}\right], & P_{\mathcal{C}}(\boldsymbol{x}) \in \partial\mathcal{C} \\ 1: & \frac{r}{\sqrt{1-r^2}} > \frac{1}{c}, & P_{\mathcal{C}}(\boldsymbol{x}) \in \mathcal{C}^{\circ} \end{cases} \tag{17}$$

where $\mathcal{C}^{\circ}$ is the interior of $\mathcal{C}$,

$$\begin{aligned} \gamma &\triangleq \rho_{\infty} ||\boldsymbol{x}|| \left(r + c\sqrt{1-r^2}\right), \\ c &\triangleq \sqrt{\frac{1-\rho_{\infty}^2}{\rho_{\infty}^2}}, \end{aligned} \tag{18}$$

and r abbreviates $r(\boldsymbol{x})$ (Equation 3). The PDF $f_S(s; \mathcal{H}_k)$ under hypothesis $\mathcal{H}_k$ for cone statistic $s(\boldsymbol{x})$ is also conditioned on the nonlinear projection $P_{\mathcal{C}}(\boldsymbol{x})$ of data stream $\boldsymbol{x}$ onto $\mathcal{C}$. We express $f_S(s; \mathcal{H}_k)$ in Appendix C (Equation C.22) using the law of total probability, along with the correlation PDF for $r(\boldsymbol{x})$ (Equation 6) to show that:

$$f_S(s; \mathcal{H}_k) \propto f_R(r(s); \mathcal{H}_k) \cdot \left|\frac{dr(s)}{ds}\right| \tag{19}$$

where r is a function of s:

$$r(s) = \rho_{\infty} s - \sqrt{1-\rho_{\infty}^2}\sqrt{1-s^2}, \tag{20}$$

and the proportionality constant for $f_S(s; \mathcal{H}_k)$ normalizes its integral over $[0, 1]$ to one. Since $f_S(s; \mathcal{H}_k)$ is functionally dependent on $f_R(r; \mathcal{H}_k)$, it is also parameterized by λ and $\lambda^{\perp}$ (Equation A.3).

To select a threshold η for event declaration, we compute the inverse, right-tail probability of $s(\boldsymbol{x})$, under $\mathcal{H}_0$, at fixed Pr_{FA}, using the Neyman-Pearson criteria:

$$\mathrm{Pr}_{FA} = \int_{\eta}^{\infty} f_S(s; \mathcal{H}_0)\, ds. \tag{21}$$

Likewise, the probability of detecting a target waveform included in $\mathcal{C}$ is the integral of the alternative hypothesis PDF over the same acceptance interval:

$$\Pr_D = \int_{\eta}^{\infty} f_S\left(s; \mathcal{H}_1\right) ds. \tag{22}$$

The comparison between Equation 21 and Equation 22 quantifies the performance of $s\left(\boldsymbol{x}\right)$ as a detection statistic for target waveforms imperfectly correlated with the template ($\rho_\infty < 1$). Fortunately, both $s\left(\boldsymbol{x}\right)$ and $f_S\left(s; \mathcal{H}_1\right)$ are functionally dependent upon the correlation statistic $r\left(\boldsymbol{x}\right)$. Consequently, both the cone statistic and its PDF are evaluable by first running a correlation detector (see Correlation Detectors) and then performing a variable transformation on the resulting statistic. Our cone detector therefore requires no additional signal scanning routines to correct for presence of template-decorrelated target signals.

Detection Probability versus Magnitude

We now quantify how decreasing values of ρ_∞ reduce the probability of detecting an explosion with a given seismic magnitude. We express this performance loss using the relative magnitude Δm between (1) a seismic source that triggers a template waveform and (2) a hypothetical seismic source producing a target waveform. In Appendix D, we show that Δm is related to the noncentrality parameter λ that shapes the cone statistic's PDF $f_S\left(s; \mathcal{H}_1\right)$. Specifically, the relative magnitude estimate between an explosion that generates a noisy waveform $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ and another explosion of magnitude m_{T} that generates a template waveform $\boldsymbol{w}$ is (Equation D.6):

$$\lambda\left(\Delta m\right) = \rho_\infty^2 \frac{\|\boldsymbol{w}\|^2}{\sigma^2} 10^{2 \cdot \Delta m}. \tag{23}$$

Here, ρ_∞ is the deterministic waveform correlation between the template and target waveform (last factor in Equation 11) and σ^2 is the variance within the data stream $\boldsymbol{x}$. The probability $\Pr_D$ that a convex cone detector with statistic $s\left(\boldsymbol{x}\right)$ identifies an explosion of magnitude $m_b = m_{\mathrm{T}} + \Delta m$ is then:

$$\Pr_D|\Delta m = \int_{\eta}^{\infty} f_S\left(s|\lambda\left(\Delta m\right);\mathcal{H}_1\right) ds, \tag{24}$$

where η is consistent with a false-alarm-on-noise probability $\Pr_{FA}$ and $|\lambda\left(\Delta m\right)$ indicates parameterization by a given relative magnitude Δm. In particular, Equation 24 quantifies the cone detector's ability to monitor for template-dissimilar waveforms ($\rho_\infty < 1$) produced by explosions of lower magnitude ($\Delta m < 0$). In the special case that target data contain an amplitude-scaled copy of the template waveform (generated by a template-colocated source), $\rho_\infty = 1$, and the cone and correlation detector are equivalent. Then Equation 24 reduces to Equation 5 and maintains the conditioning on Δm. In all other cases, the difference in detection power (λ) between the cone and correlation detector is quadratic in ρ_∞ and $\Pr_D|\Delta m$ thereby defines a decreasing function of ρ_∞ for a fixed relative magnitude.

Application: Explosions in North Korea

We consider a CTBTO monitoring mission to detect explosions at the North Korean Nuclear Test Site (NKTS) using a correlation and cone detector with a sub-network of IMS receivers. Our goal is to bound the probability of detecting an explosion-triggered waveform originating from the NKTS, using a template with a dissimilar shape, where the target explosion is smaller in magnitude than that of previous tests. To pursue this goal, we first process target data recording the previous test explosions using a maximum likelihood multichannel correlation detector (Equation 3). We then use the results of this routine to parameterize a competing cone detector (Equation 16) that we also apply to our target data. Based on the results of these routines, we assess our cone detectors' capability to detect explosions at increasing values of waveform dissimilarity ($1 - \rho_\infty$) and over a range of prescribed magnitudes. Last, we use a semi-empirical test to evaluate the performance of these competing detectors over range of source sizes and compute the probability of detecting of small explosions at the 2006 test site with a template recording the 2013 waveforms.

Cone Detector Demonstration

Our data include three explosions: the 09-Oct-2006, 25-May-2009 and 12-Feb-2013 announced nuclear tests, each conducted at the NKTS and separated by $\leq$ 2.5 km (Begnaud et al. (2011)) (Table 1). All three events were recorded on IMS seismic arrays deployed in Japan (MJAR) and Kazakhstan (MKAR) (Figure 3). Additional arrays also recorded the first explosion, and several closer IMS arrays certified after 2007 recorded the later tests. While this study uses MJAR and MKAR exclusively, extension to additional NKTS proximal arrays is straightforward.

To proceed, we collected data from MJAR and MKAR and prepared them for correlation processing by detrending each time series in 1 hr segments, demeaning these data with a 60 s running average, and then bandpass filtering the results over 1.5-7.5 Hz using a four-pole Butterworth filter. We then scanned the 2006 and 2009 explosion data using an identically pre-processed template waveform $\boldsymbol{w}$ that we manually extracted from 2013 explosion records by picking high-SNR P-phase segments on a subset of channels (Figure 4). To estimate a threshold for match detection ($\hat{\eta}$), we established a 10^{-8} false-alarm-on-noise probability and inverted for the lower integration bound of Equation 4:

$$\hat{\eta} = \underset{\eta}{\operatorname{argmin}} \left\{ \left| 10^{-8} - \int_{\eta}^{1} f_R \left(r | \hat{N}_E ; \mathcal{H}_0 \right) dr \right| \right\} \tag{25}$$

Here, $|\hat{N}_E$ indicates that the null PDF for r is conditioned on an estimate of the effective degrees of freedom N_E in data $\boldsymbol{x}$ (Equation B.1), and $\hat{\eta}$ is therefore estimated from an imperfectly known density function; Carmichael and Hartse ((2016)) describe the additional operational details of this detector. Given these estimates, our correlation routine identified both the 2006 and 2009 events, but with unequal significance. Specifically, the peak statistic recording the 2006 event measured about half the peak statistic recording the 2009 event ($r(\boldsymbol{x}) \approx 0.44$ versus $r(\boldsymbol{x}) \approx 0.86$). Hypocentral locations reported in the Reviewed Event Bulletin (REB; Table 1) also indicate that the 2006 event was further separated from the 2013 source when compared to the more proximal 2009 nuclear test. This spatial separation almost certainly reduced the template/target waveform similarity. To quantify such deterministic decorrelation, we applied Equation 12 and estimated ρ_∞ between the template and each target explosion. These estimates each exceeded the peak correlation statistic $r(\boldsymbol{x})$ and respectively gave $\hat{\rho}_\infty = 0.5$ (2013 versus 2006) and $\hat{\rho}_\infty = 0.88$ (2013 versus 2009). We then used the first estimate to reprocess the explosion records with a convex cone detector parameterized by $\rho_\infty = 0.5$. These cone detection routines identified both the 2006 and 2009 explosion waveforms, but now with almost equal significance. The histograms for these test statistics $s(\boldsymbol{x})$ matched their theoretical PDFs within less than 4% relative error and gave confidence to our threshold estimates and support for the cone detector theory developed in Appendix C.

Having demonstrated the capability of the cone detector to identify explosion-triggered waveforms using IMS stations, we employed the 2013 template $(\boldsymbol{w})$ to estimate the probability of detecting other waveforms, triggered by smaller explosions. We focused on waveforms showing template cross-correlation values comparable to those measured from the 2009 and 2006 tests. This process included two assessments. Our first assessment analyzed the performance of the cone detector over a range of deterministic correlation values for a fixed, reference magnitude. We selected this magnitude from a prior IMS capability study that concluded $m_b = 3.25$ events originating from North Korea had a 0.90 probability of being detected on three or more stations with a power detector (Kværna and Ringdal (2013)) (with IMS coverage in 2013). We then used this reference event to estimate the conditional probability of detecting waveforms triggered by a similarly sized explosion with our cone detector, when the target waveforms were decorrelated with our template. We considered particularl values of ρ_∞ consistent with the two earlier explosions (2006, 2009). Our second assessment estimated cone detection probabilities at fixed ρ_∞ and variable explosion magnitude. Specifically, we calculated the probability of detecting waveforms triggered by explosions collocated with the 2006 and 2009 tests over a range of prescribed magnitudes. To do so, we applied cone detectors parameterized by our two estimates of ρ_∞ and thereby constructed two distinct detection probability curves. In both of these assessments, we computed the relevant noise-statistics ($\hat{\sigma}$, $\hat{N}_E$ and $\hat{\eta}$) that shaped these curves using data recorded 1800 s before and 3600 s after each explosion.

Semi-Empirical Performance Comparison

Our capability assessments (see Cone Detector Demonstration) were limited by data measured immediately before and after the 2006 and 2009 tests. To extend our analyses, we processed records that included waveforms like those measured during the 2006 explosion over a range of magnitudes and noise conditions. We constructed these waveforms by contaminating scaled copies of the 2006 test records with additional noise recorded by MKAR and MJAR during the previous day (08-Oct-2006). We then used these data to model waveforms triggered by explosions ranging from $m_{\mathrm{b}} = 2.9$ to $m_{\mathrm{b}} = 4.3$ units in magnitude and processed them with correlation and cone detectors. This procedure included four stages. First, we scaled the 2006 test records by $10^{\Delta m}$ to mimic waveform amplitudes produced by a hypothetical explosion of template-relative magnitude Δm. Second, we added these scaled waveforms ≈ 250 consecutive times to 24 adjacent, 1 hr noise windows recorded over 08-Oct-2006. We did not remove any background seismicity, identified by a power detector, or otherwise. Third, we processed these degraded data with a correlation and cone detector that each included the template recording the 2013 test (Figure 5). The cone detector operated with $\rho_\infty = 0.5$ and the same false-alarm-on-noise constraint as the correlation detector (Equation 21). Last, we counted the total number of true detections and normalized this count by the total number of expected detections in each magnitude bin; we discounted false detections from background seismicity. This produced two empirical receiver-operating characteristic (ROC) curves that measured the probability of detecting an explosion at the 2006 test location with the 2013 explosion template. After computing these ROC curves, we calculated the theoretical performance of each detector over a 100-point magnitude grid for comparison (Equation 24). To do so, we first estimated the parameters σ^2, λ and N_E that shaped these curves. To estimate σ^2, we removed potentially biasing signal from each data stream channel with a power (STA/LTA) detector and computed hourly averages of the variance from the remaining data. We then summed these channel-wise variance estimates to compute the multichannel variance for each hour of 08-Oct-2006. We similarly computed hourly estimates of λ (Equation 23) from our preceding variance estimates and prescribed value of $\rho_\infty = 0.5$. This produced 24, hour-specific ROC curves that measured the predicted performance of each detector over our 100 point magnitude grid.

Results

The correlation and cone detectors each identified the 2009 and 2006 explosions at MJAR and MKAR with the template recorded from the 2013 test (Figure 6). These detectors produced statistics with several dissimilarities, however. The correlation statistics measured a zero mean in the presence of noise and a difference of 0.4 correlation units between target detections (Figures 6a and 6c). The cone statistics measured a non-zero mean (0.88) in the presence of noise and nearly identical, unit-valued target detections (Figures 6b, 6d and 6e). The inter-event similarity in the sample mean of the cone statistic and the near-agreement in its detection maxima follows from the (near) inclusion of both target waveforms in the same convex set $\mathcal{C}$. Specifically, the high mean of the statistic is a consequence of the cone geometry: Gaussian noise produces larger projections onto the cone boundary than onto the cone vertex (template). The near agreement in peak value is a consequence of the signal model: both target waveforms populate $\mathcal{C}$ and their noisy realizations therefore produce similarly sized projections onto it (Equation 17). These detector differences indicate that the cone statistic is not interpretable with intuition developed from correlation detector processing.

Cone Detector Parameterization

Our first assessment quantifies the decrease in waveform detection probability versus template similarity for explosions of fixed magnitude near the NKTS. This assessment suggests that little performance is lost for deterministic correlation values ρ_∞ greater than 0.85 (Figure 7) when the target source is a small explosion ($m_b = 3.25$). Specifically, the probability of detecting waveforms triggered by such an $m_b = 3.25$ explosion, which deterministically cross-correlate with our template at $\rho_\infty = 0.88$ (like the 2009 test), is $\Pr_D \approx 0.98$. Our cone detector would likely have missed such a low magnitude source if it produced waveforms as poorly correlated with the template as those recorded during the 2006 test. In this second case, the diminished signal semblance ($\rho_\infty \approx 0.5$) resulted in a detection probability below 20% for the same false-alarm-on-noise probability.

Our second assessment shows that explosions near the NKTS that trigger template-dissimilar waveforms must be substantially larger in magnitude than those that trigger more template-similar waveforms to achieve the same detection probability (Figure 8). For specificity, we compare waveforms that deterministically correlate with our template as well as those produced by 2006 and 2009 tests and have the same 0.90 probability of being detected. In this case, an $m_b = 3.42$ explosion at the 2006 source location has the same chance as being identified with our template as an $m_b = 3.18$ explosion at the 2009 source location (Figure 8, white markers). Such a hypothetical $m_b = 3.18$ explosion is effectively undetectable with the cone detector ($\mathrm{Pr}_D < 6\%$) if it produces waveforms as poorly correlated with the template as those produced by the 2006 explosion. Other explosions located near the 2006 test hypocenter would likely produce such template-dissimilar signals. Regardless of the cause of such waveform dissimilarity, our example suggests that ≈ 0.4 units of cross-correlation reduction can "cost" 0.25 magnitude units of detection capability when using certain arrays within the IMS to monitor the NKTS.

Semi-Empirical Comparison with the Correlation Detector

Figure 9a shows that our empirical correlation results deviate from predicted detection capability that we derived from the hypothesis test of Equation 2. Specifically, these empirical results underperform relative to the theoretical curve that models the target waveform as an amplitude-scaled copy of the template buried in noise. In fact, the time-averaged empirical ROC curve (solid line) falls outside the expected range of predictions (shaded area) for any detection probabilities that exceed about 0.4. The empirical detection results from the cone, in contrast, agree well within our associated range of predictions. In this latter case, the time-averaged empirical ROC curve (Figure 9b, solid line) is interior to the shaded region that marks our predictions, for all but the highest magnitude values. Importantly, this agreement holds at the 0.90 detection probability threshold that is often used as a monitoring benchmark (Gibbons et al. (2012); Kværna and Ringdal (2013)).

The width of each shaded region in Figure 9 also indicates that the detector performance can range markedly by time of day. Specifically, the detection probability of quasi-repeatable explosions with fixed magnitude spanned almost 0.7 probability units over 24 hours (Figure 9, error bars), with hour 1 (UCT time) giving the lowest routine-averaged detection probability for both detectors. Our estimate for the effective degrees of freedom (estimated as $\hat{N}_E$) also show a corresponding minimum at this hour. These minima indicate a relatively strong correlation structure present within the data, which may be attributable to microseismic noise. Our manual inspection of waveform data recorded at MJAR during this time showed coincident, a high amplitude ≤ 2 Hz spectral peak (not plotted). The presence of this peak, and it's relative absence during other times, suggests that narrowband interference is potentially responsible for the degradation of our detector's capability. We pursue the influence of narrowband interference on detection capability in Part II of this work and consider other sources of error to the present study in Appendix E.

Cone versus Correlation Detector Performance

The empirically-derived ROC curve for the correlation detector indicates a marginal outperformance relative to that of the cone detector (Figure 10, stair plots). We expected this performance gap, since the cone detector requires a maximum likelihood estimate for an N-dimensional signal $\boldsymbol{u}$ and therefore includes additional parametric uncertainty. The correlation detector, by comparison, requires estimation of a scalar maximum likelihood amplitude A. Despite this disparity, the performance gap between the empirical versus theoretical ROC curve for the correlation detector is substantial larger than that between the empirical cone-versus-correlation detector ROC curves. Figure 10 shows a solid curve that marks the lower-bound on the shaded region from Figure 9a and measures the minimum expected discrepancy. The mismatch between the empirical and theoretical curves indicate that $m_{\mathrm{b}} \approx 3.6$ explosions separated ≈ 2 km from the template source (the 2013-to-2006 source separation distance) have a $> 16\%$ smaller chance of being detected than predicted in the same noise conditions (vertical lines). Because the empirically-derived ROC curve for the cone detector is bounded within the range of predictions over the same magnitude interval, it shows no analogous discrepancy. Therefore, the cone detector shows better predictive capability compared to observations and only marginally lower performance.

False Alarm Comparison

In addition to accommodating template-dissimilar waveforms, our cone detector also returned 23% fewer nuisance alarms when compared with our correlation detector (Figure 11). We did not anticipate this fortunate reduction, since both detectors operated at the same 10^{-8} false-alarm-on-noise probability constraint. We explain this nuisance alarm reduction geometrically: admission of any deterministic template/target waveform dissimilarity requires a comparatively rapid inflation of detector thresholds to maintain a false-alarm on noise probability consistent with that of a reference correlation detector. This inflation occurs because the vector-space representations of noise $\boldsymbol{n}$ have greater probability of producing substantial projections onto the target cone boundary ($\partial\mathcal{C}$) than they do onto the one-dimensional subspace span($\boldsymbol{w}$) representing the waveform template (the cone vertex). Conic projections of noise are also realized more often during processing (scanning), since all vectorial orientations of noise are equally likely. Non-target waveforms ($\notin \mathcal{C}$) similarly have a greater chance of giving substantial projection onto the cone, again because these signals are geometrically closer to the cone boundary than to the cone vertex. However, a fixed false-alarm-on-noise rate $\Pr_{FA}$, identical to that of the correlation detector, requires elevated event declaration thresholds for the cone case so that noise-projections are constant. As a result, any deterministic correlation loss between target and template waveforms requires a comparatively large, nonlinear increase in cone detector thresholds (see Appendix C).

Discussion

Our theoretical and observational work both indicate that unanticipated dissimilarities between explosively triggered waveforms can contradict the expected monitoring performance of waveform correlation detectors that exploit IMS data (Figure 9a, Figure 10). This waveform discrepancy is particularly important when monitoring for small explosions that may be located several km from prior test locations. Unfortunately, correlation detectors do not correctly accommodate such uncertainty. While the probabilistic expression of these detectors (the PDFs) can be parameterized to include deterministic correlation losses (Carmichael and Hartse (2016)), this modification contradicts the correlation detector's operational model. Namely, that target waveforms have the same shape as the template, and that their relative magnitudes are related through a log-ratio of amplitudes. Observations from explosions at the NKTS contradict this simplistic assumption. Waveforms generated by the spatially separated explosions located there instead exhibit destructive inter-event interference and are therefore geometrically dissimilar (not proportional in amplitude).

The convex cone detector that we have introduced in this study addresses several of these monitoring challenges. Like a correlation detector, it searches for waveforms that are similar to the template waveform in shape. Unlike a correlation detector, it admits a deterministic level of uncertainty in matching this template's shape, but maintains the same false alarm on noise probability. This admission of waveform uncertainty expands the set of target signals from those that are only amplitude-scaled copies of the template to all waveforms that significantly correlate with the template's underlying signal (as measured by ρ_∞). The associated preservation of a false-alarm on noise probability means that a more inclusive signal set does not imply more non-target detections. The size of this set can also be tuned to any size by adjusting parameter ρ_∞. Smaller values of ρ_∞ then give a larger signal space and correspond to greater uncertainty in the template/target waveform similarity.

There is an unavoidable tradeoff between target signal uncertainty and detection power, however. That is, decreased similarity in template/target waveform geometry penalizes the cone detector's performance relative to that of the correlation detector (Figure 10). We suggest that this cost is offset by three attributes: (1) a more physically representative signal-present hypothesis that provides (2) an improved predictive capability over that of the correlation detector (Figure 9b), and (3) a reduction in false alarms on background seismicity (Figure 11). The operational cost of including a cone detector in processing pipelines that include correlation detectors is also negligible.

This is because the cone statistic requires no additional scanning routines. It only requires a functional transformation of the correlation statistic and recalculation of the detection threshold used in the decision rule. This transformation, however, is also conditional on the value of the correlation statistic because the cone projection is nonlinear (Equation 17). Consequently, the mapping from $r(\boldsymbol{x})$ to $s(\boldsymbol{x})$ includes implementing a decision tree (piecewise function). This also includes little computational cost in pipelines that implement codes with vector arithmetic enabled. We concede, nevertheless, that our performance comparison is very limited and requires additional empirical study with a larger range of sources, magnitudes, and standoff distances that we have included here.

Despite our emphasis on the limitation of correlation detectors, they remain a valuable seismic monitoring tool. This is because they are optimal detectors for target waveforms buried in normally distributed noise with known shapes but unknown amplitudes. Their implementation requires scrutiny, however. While correlation detectors can identify very low SNR waveforms, the sources generating such signals cannot generally be identified on the basis of detection alone. We suggest that such identification requires that template and target waveforms share a large deterministic cross-correlation value (Equation 11) in the absence of other seismic evidence for source identity. In such cases, our cone detector is a useful tool to assess ostensible uncertainty in the application of correlation detectors in post-processing stages, after initial detections are made. Such uncertainty analysis is crucial for monitoring small magnitude sources at near-regional to local standoff distances, where additional data on source type may be unavailable (e.g., Zhang and Wen (2015)). Our analysis, however, has so far only treated far-regional and teleseismic IMS array measurements. We selected these instruments because (1) they recorded all three of the announced explosions conducted before 2016 and (2) the relatively low SNR values for the 2006 test is demonstrative of expected SNR values from smaller explosions recorded at a closer standoff distance. We therefore emphasize that our methodology is equally applicable, and probably more important at these closer monitoring distances.

Finally, while our study focused on waveforms triggered by spatially separated explosions, alternative physical processes involving other seismic sources could also induce deterministic waveform differences. Other causes of decorrelation include non overlapping corner frequencies of the source time function (Ford and Walter (2015)), damage production near the source (Patton and Taylor (2008)), or coincident tectonic release. However, such processes are difficult to observe, whereas relative hypocentral locations at the NKTS are generally known within $\approx$ 200 m or less. Any additional study of our cone detector will likely require certain ground-truth locations to assess influence from such alternative physical mechanisms that induce waveform changes.

Conclusion

Waveform correlation over predicts the capability of IMS arrays to detect non-repeatable low magnitude explosions and can underestimate the associated relative magnitude of any correctly identified explosions. Such assessment errors comprise a continuing verification challenge for the CTBTO, especially for monitoring smaller magnitude tests. These challenges require detectors that are explicitly derived to accommodate waveform dissimilarity. We have provided such a detector. This "cone" detector accommodates deterministic differences between template and target waveforms and can detect decorrelated signals without increasing false alarms on noise. In addition, it only requires records from one previous explosion, unlike a higher-rank subspace detector that needs measurements recording several distinct events. While more study is required, we recommend implementing cone detector prototypes in processing pipelines to more effectively monitor for low-magnitude explosions that may be spatially distributed over a test site.

To our knowledge, this work is the first to probabilistically quantify how deterministic uncertainty in template/target waveform similarity adds to noise to comprehensively degrade waveform correlation detector performance. This degradation and analysis has consequences for geophysical monitoring operations that are unrelated to explosion detection as well. Namely, these detectors are applicable in detection problems where inter-event waveform evolution is driven by spatiotemporal complexity, such as in developing aftershock sequences, or in icequake detection, where the source medium moves relative to the receiver.

We concede that this work did not address several research challenges of waveform correlation. Among these is nuisance seismicity, which represents an error in the signal absent model. Our future work (Part II) includes addressing these errors and thereby modifying correlation detectors to accommodate template-correlated background seismicity into our hypothesis tests.

Data and Resources

We acquired seismic data recorded at MJAR and MKAR on 03-Jan-2016 from the International Monitoring System (IMS), which is available from the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) in Vienna, Austria. We processed all data with the `CORAL` toolbox written in `Matlab` by Ken Creager, and later expanded by Joshua D. Carmichael while at the University of Washington and Los Alamos National Laboratory.

Acknowledgment

Mike Cleveland, Hans Hartse and Char Rowe improved emphasis and writing clarity. Amanda Ziemann provided her expertise regarding detection theory. Development of initial detection theory work for icequake analyses was supported by NSF award 1023382. The bulk of the seismic work and application was supported by the Office of Nuclear Detonation Detection, NA-222. Los Alamos National Laboratory is operated for the Department of Energy by Los Alamos National Security, LLC, under Contract DE-AC52-06NA25396.

References

Begnaud, M. L., Phillips, W. S., Stead, R. J., and Pabian, F. V. (2011). Constraints on the Locations of the 2009 and 2006 DPRK P'unggyei Explosions from Fusion of Improved Relative Relocation and Geospatial Analysis. *Defense Research Review*, 19(1):1–16.

Bobrov, D., Kitov, I., and Zerbo, L. (2014). Perspectives of Cross-Correlation in Seismic Monitoring at the International Data Centre. *Pure and Applied Geophysics*, 171(3-5):439–468.

Carmichael, J. D. (2013). *Melt-Triggered Seismic Response in Hydraulically-Active Polar Ice: Observations and Methods*. PhD thesis, University of Washington.

Carmichael, J. D. and Hartse, H. (2016). Threshold Magnitudes for a Multichannel Correlation Detector in Background Seismicity. *Bulletin of the Seismological Society of America*, 106(2):18.

Carmichael, J. D., Nemzek, R., Arrowsmith, S., and Sentz, K. (2016). Fusing Geophysical Signatures of Locally Recorded Surface Explosions to Improve Blast Detection. *Geophysical Journal International*, 204(3):1838–1842.

Dodge, D. and Walter, W. (2015). Initial Global Seismic Cross-Correlation Results: Implications for Empirical Signal Detectors. *The Bulletin of the Seismological Society of America*, 105:240–256.

Ford, S. R. and Walter, W. R. (2015). International Monitoring System Correlation Detection at the North Korean Nuclear Test Site at Punggye-'ri with Insights from the Source Physics Experiment. *Seismological Research Letters*, 86(4):1160–1170.

Gibbons, S. J. and Ringdal, F. (2006). The Detection of Low Magnitude Seismic Events using Array-Based Waveform Correlation. *Geophysical Journal International*, 165(1):149–166. PT: J; TC: 49; UT: WOS:000239689500013.

Gibbons, S. J., Ringdal, F., and Kvaerna, T. (2012). Ratio-to-moving-average seismograms: a strategy for improving correlation detector performance. *Geophysical Journal International*, 190(1):511–521. PT: J; NR: 48; TC: 0; J9: GEOPHYS J INT; PG: 11; GA: 956HO; UT: WOS:000305080900037.

Harris, D. B. (1989). Characterizing Source Regions with Signal Subspace Methods: Theory and Computational Methods. Technical Report UCID-21848, Lawrence Livermore National Laboratory.

Harris, D. B. (1991). A Waveform Correlation Method for Identifying Quarry Explosions. *Bulletin of the Seismological Society of America*, 81(6):2395–2418.

Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice-Hall PTR, Englewood Cliffs, N.J. ID: 26504848.

Kay, S. M. (1998). *Fundamentals of Statistical Signal Processing: Detection Theory*. Prentice-Hall Inc., Upper Saddle River, New Jersey, USA, 1st edition.

Kværna, T. and Ringdal, F. (2013). Detection Capability of the Seismic Network of the International Monitoring System for the Comprehensive Nuclear-Test-Ban Treaty. *Bulletin of the Seismological Society of America*, 103(2A):759–772.

MacCarthy, J. K., Hartse, H., Greene, M., and Rowe, C. (2008). Using Waveform Cross-Correlation and Satellite Imagery to Identify Repeating Mine Blasts in Eastern Kazakhstan. *Seismological Research Letters*, 79(3):393–399.

Moreau, J.-J. (1962). Décomposition Orthogonale dun Espace Hilbertien Selon deux cônes Mutuellement Polaires. *CR Acad. Sci. Paris*, 255:238–240.

Patton, H. J. and Taylor, S. R. (2008). Effects of Shock-Induced Tensile Failure on mb-ms Discrimination: Contrasts between Historic Nuclear Explosions and the North Korean Test of 9 October 2006. *Geophysical Research Letters*, 35(14).

Schaff, D. (2010). Improvements to Detection Capability by Cross-Correlating for Similar Events: a Case Study of the 1999 Xiuyan, China, Sequence and Synthetic Sensitivity Tests. *Geophysical Journal International*, 180(2):829–846.

Scharf, L. L. and Friedlander, B. (1994). Matched Subspace Detectors. *IEEE Transactions on Signal Processing*, 42(8):2146–2156.

Stark, H. and Yang, Y. (1998). *Vector Space Projections: A Numerical Approach to Signal and Image Processing, Neural Nets, and Optics*. Wiley John and Sons, Incorporated, New York, NY.

Weichecki-Vergara, S., Gray, H. L., and Woodward, W. A. (2001). Statistical Development in Support of CTBT Monitoring. Technical Report DTRA-TR-00-22, Southern Methodist University.

Zhang, M. and Wen, L. (2015). Seismological Evidence for a Low-Yield Nuclear Test on 12 May 2010 in North Korea. *Seismological Research Letters*, 86(1):138–145.

List of Tables

List of Figures

Table 1: Reviewed Event Bulletin and Correlation Data for North Korean Tests

Origin Time	Loc.	m_b	Peak CC Stat.	Peak Cone Stat.
10/09/2006 (282) 01:35:27.8	41.28546, 129.10878	4.2	0.44	0.98
05/25/2009 (145) 00:54:43.0	41.29144, 129.08307	4.6	0.86	1
02/12/2013 (043) 02:57:51.1	41.28853, 129.08142	5.0	1	1

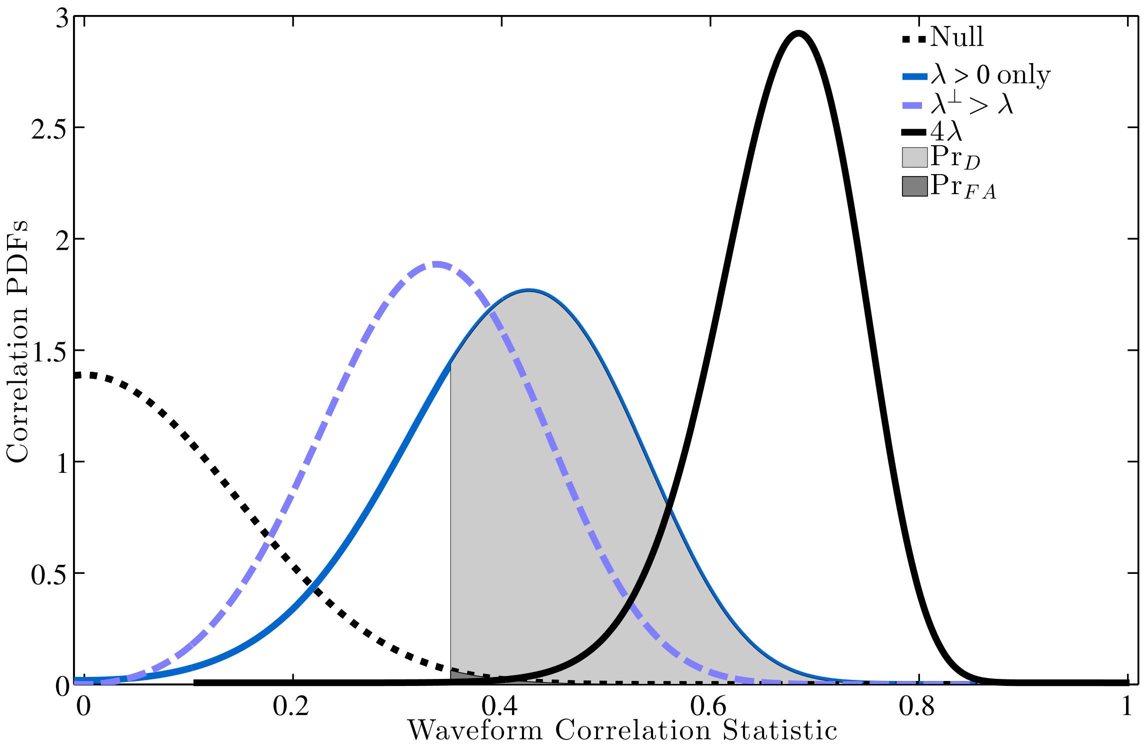

Figure 1: Notional PDF curves for correlation statistics $r(\boldsymbol{x})$ when $N_E = 50$, for several noncentrality parameters. The left-most dashed curve shows the density for $r(\boldsymbol{x})$ in the noise-only case ($\lambda = \lambda^{\perp} = 0$). The leftmost solid curve shows the density for $r(\boldsymbol{x})$ when a target signal is present ($\lambda = 10$, $\lambda^{\perp} = 0$) The rightmost dashed curve shows the density for $r(\boldsymbol{x})$ when a partially correlated signal is present ($\lambda = 5$, $\lambda^{\perp} = 8.66$). The rightmost solid curve shows the density when $\lambda = 40$. The vertical line indicates threshold $\eta = 0.35$; the shaded area beneath the null distribution where $r(\boldsymbol{x}) > \eta$ shows the false alarm-on-noise rate, while shading beneath the solid curve shows the probability of detecting a target waveform when $\lambda = 10$.

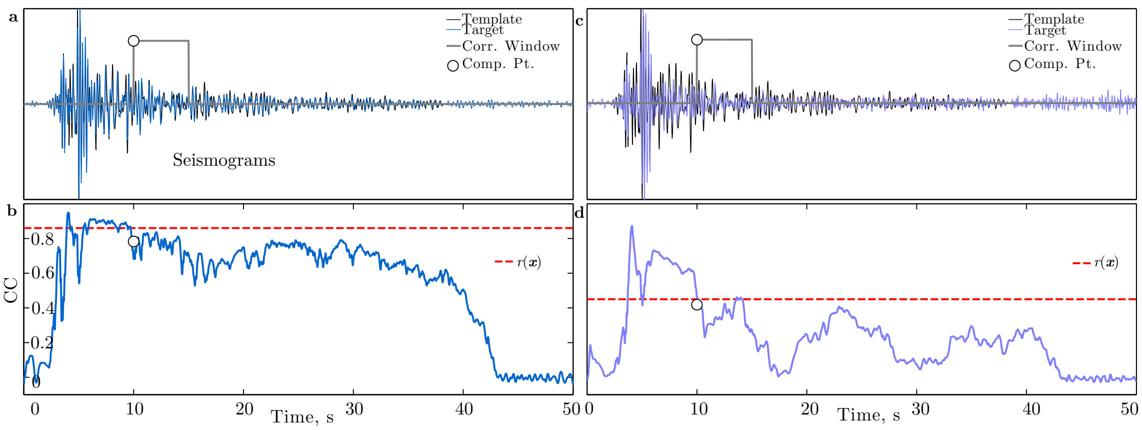

Figure 2: Two single channel examples of template/target waveform correlation with differing levels of signal similarity. The left panel shows relatively high inter-event waveform similarity, and the right window shows comparatively low waveform similarity; all waveforms are normalized to unit peak amplitude. **a**: A template waveform $\boldsymbol{w}$ (darkest seismogram) aligned with a target waveform at peak correlation to subsample precision. The square function indicates the snapshot of a 5 s cross-correlation window that scans over the pre-aligned data to compute a localized correlation coefficient. The window time stamp is indicated by the circular marker. **b**: Localized correlation coefficients (CC) computed from the 5 s moving window in **a** and indicated by the circular marker. Limited data variability suggests that a waveforms are reasonably proportional in amplitude; samples after $\sim$40 s mark where signal drops below the noise. The horizontal line shows the bulk correlation coefficient (CC) computed over the entire 50 s window. **c**: Same as **a**, but using a target waveform with less similarity with the template. **d**: Same as **b**, but including the less template-similar target waveform shown in **c**. In this case, the data variability is more pronounced than that shown in **b**. The template and target data are disproportionate in amplitude; data near 10 and 13 s are locally correlated (CC = 0.45), whereas data near 17 s are nearly decorrelated (CC $\approx$ 0).

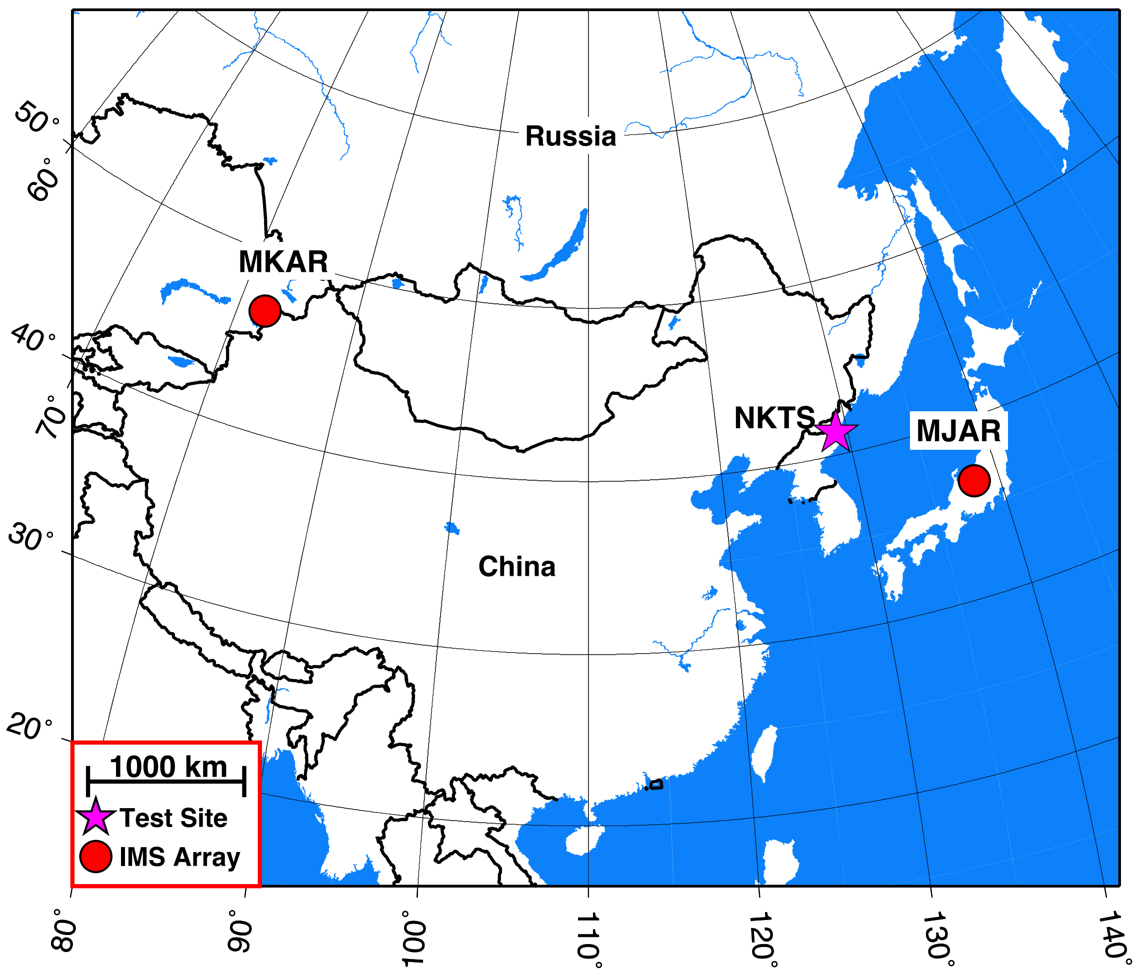

Figure 3: Geographical location of the North Korean Nuclear Test site (NKTS) and three IMS arrays. MJAR and MKAR recorded all announced nuclear test conducted by North Korea (2006, 2009, 2013, 2016).

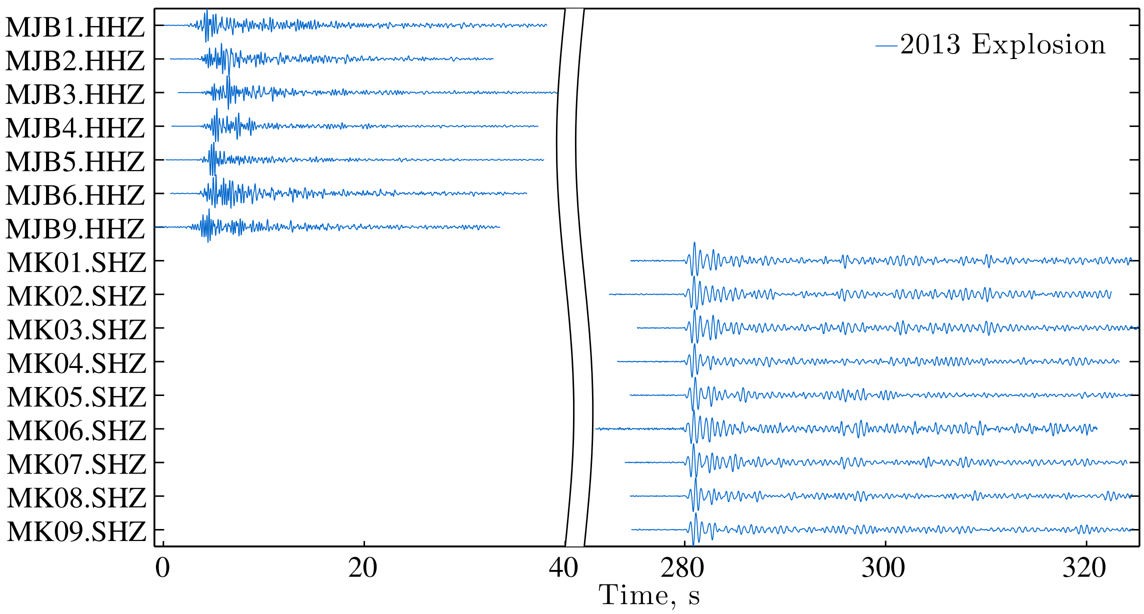

Figure 4: Short period, bandpass filtered velocity records of the 2013 North Korean test recorded at IMS arrays MJAR and MKAR.

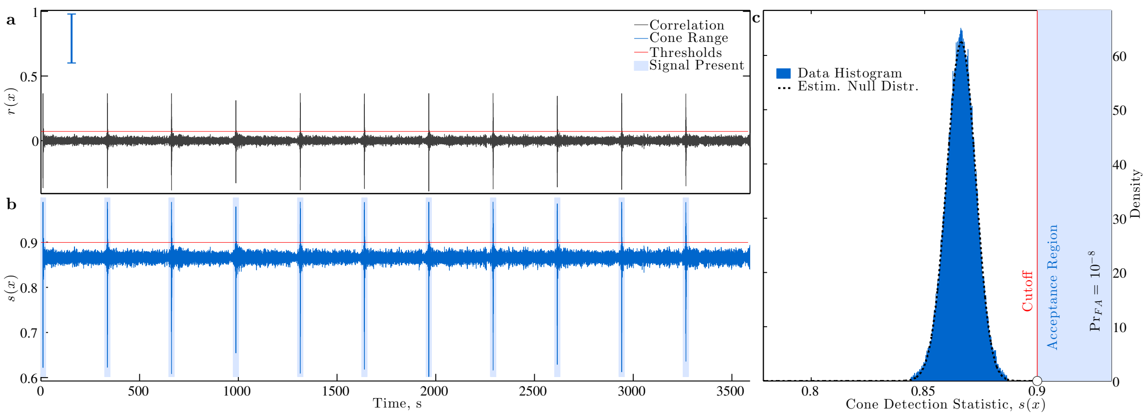

Figure 5: A comparison between semi-empirical detection statistics. Target data include a summation of (1) seismograms recorded on 08-Oct-2006 at MJAR and MKAR with (2) amplitude-scaled waveforms recording the first announced nuclear test from North Korea the next day. Template waveforms are shown in Figure 4. **a**: The correlation statistic $r(\boldsymbol{x})$ (Equation 3) computed by processing 1 hr of data. The horizontal line shows the $\mathrm{Pr}_{FA} = 10^{-8}$ threshold. Peaks in the time series indicate detections on 11 noise-contaminated waveforms. The vertical "errorbar" at the top left shows the range of corresponding cone statistic values. **b**: The waveform cone statistic $s(\boldsymbol{x})$ (Equation 17) for the 2006 target data, shown with $\mathrm{Pr}_{FA} = 10^{-8}$ detection thresholds (horizontal line). Lightly shaded data identifies waveforms. **c**: The solid histogram shows the empirical PDF computed estimated from $s(\boldsymbol{x})$ data. The dark dashed curve shows the predicted null distribution using parameter estimates from these data. The vertical line corresponds to the horizontal line at left. The shaded region indicates the acceptance region ($s(\boldsymbol{x}) > \eta$) for event detection and corresponds to the shading for $s(\boldsymbol{x})$ in **b**

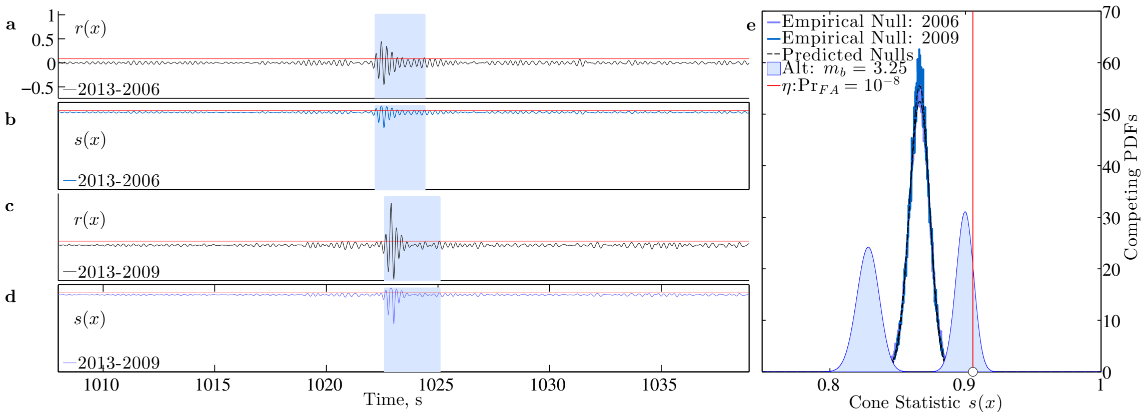

Figure 6: Various detection statistics computed from the three North Korean tests. The left panels' horizontal axis begins 900 s before REB origin times and shaded data segments show time windows centered on points where each data statistic exceeds its respective threshold. **a**: The correlation statistic $r(\boldsymbol{x})$ (Equation 3) computed by processing 2006-test data recorded at MJAR and MKAR with template waveforms extracted from the 2013 explosion. The horizontal line shows the $\Pr_{FA} = 10^{-8}$ threshold. **b**: The waveform cone statistic $s(\boldsymbol{x})$ (Equation 17) for the 2006 target data, shown with equivalent $\Pr_{FA} = 10^{-8}$ detection thresholds. **c**: Same as **a**, but for the 2009 test data. **d**: Same as **b**, but for the 2009 test data. **e**: A family of histograms and predicted (theoretical) PDFs that describe the convex cone statistics. Middle, nearly identical solid PDF curves show the histograms computed using $s(\boldsymbol{x})$ data in **b** and **d**. Nearly identical dashed curves show predicted null distributions using parameter estimates from these data. Filled PDF curves give the alternative PDF for $s(\boldsymbol{x})$, assuming noncentrality parameters consistent with a low SNR $m_b = 3.25$ event. The vertical line corresponds to the horizontal line in **b**.

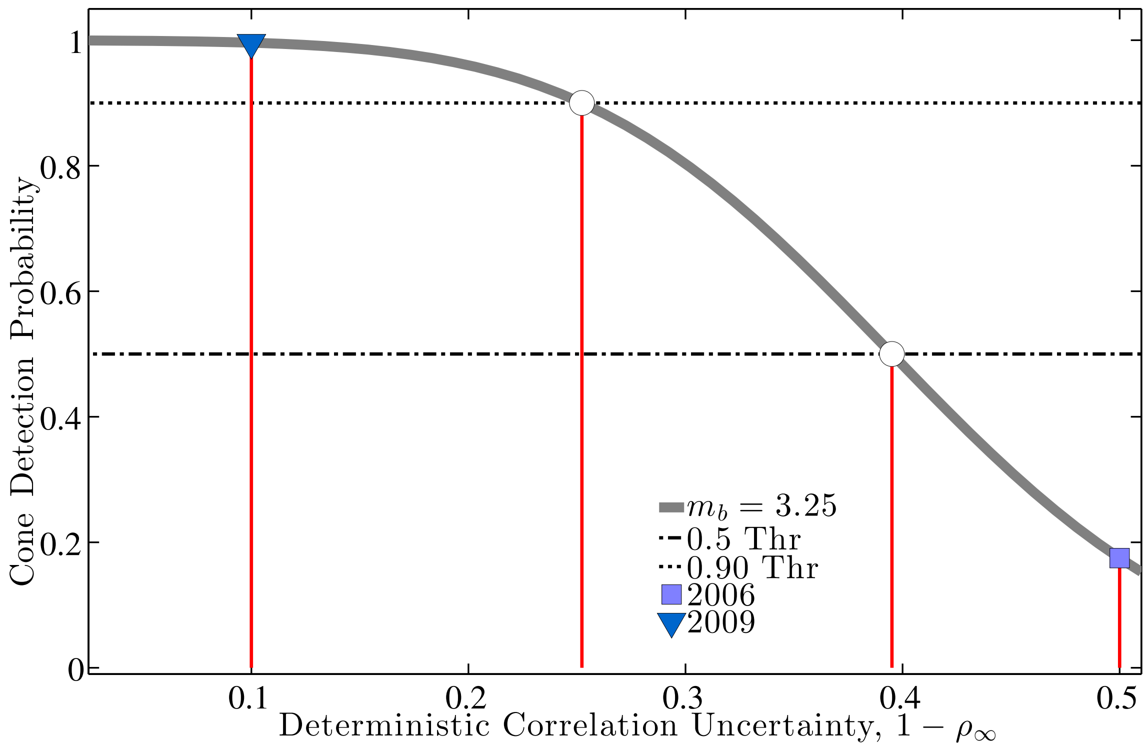

Figure 7: The performance of the cone detector versus waveform similarity $1 - \rho_\infty$ at fixed magnitude. The thick curve shows the probability of detecting an $m_b = 3.25$ explosion using a template extracted from the 2013 test (Figure 4). The triangular marker indicates approximate ρ_∞ values associated with the 2009 test. The square marker indicates the approximate ρ_∞ value associated with the 2006 test. Circular markers show where the 0.9 and 0.5 detection probability lines intersect the performance curve. Vertical lines show deterministic correlation values associated with each threshold probability.

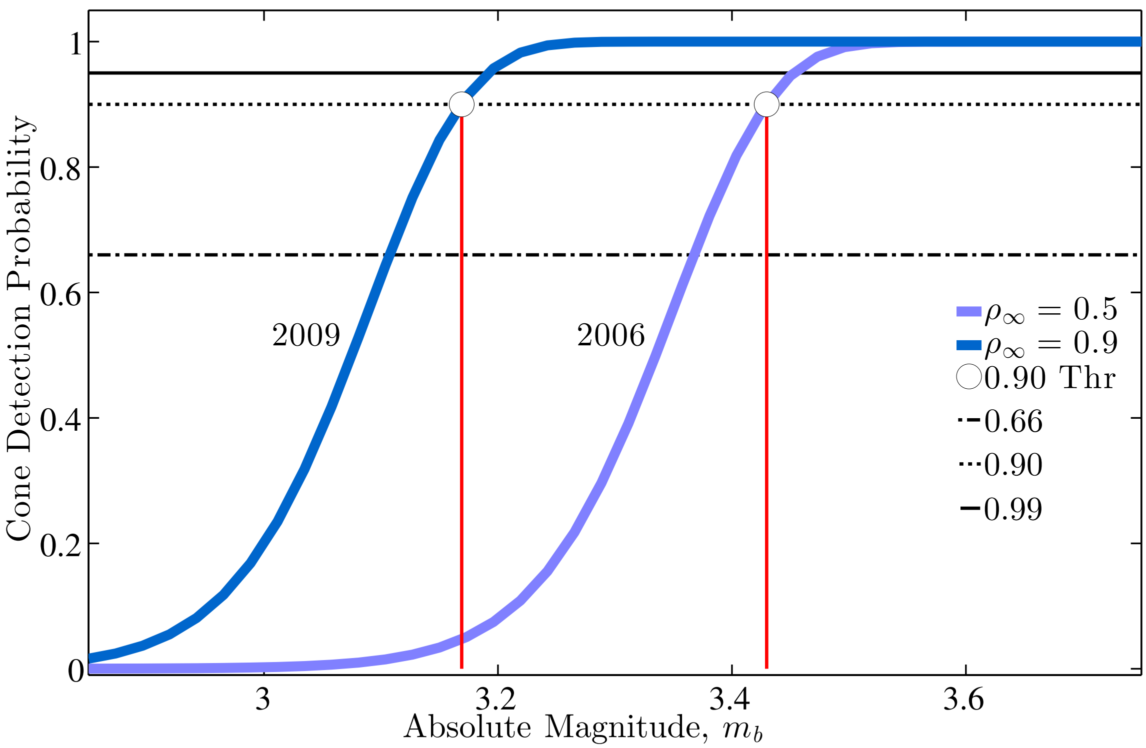

Figure 8: Detection performance of the cone detector at fixed waveform similarity ρ_∞ and variable explosion magnitude. Curves show the probability of detecting an explosion colocated with the 2009 (left) or 2006 source (right) using a template extracted from the 2013 test (Figure 4). Horizontal lines show (from top to bottom) 0.95, 0.9 and 0.66 detection probability thresholds. Circular markers show where the 0.9 detection probability lines intersect each performance curve, and vertical lines show the corresponding magnitudes.

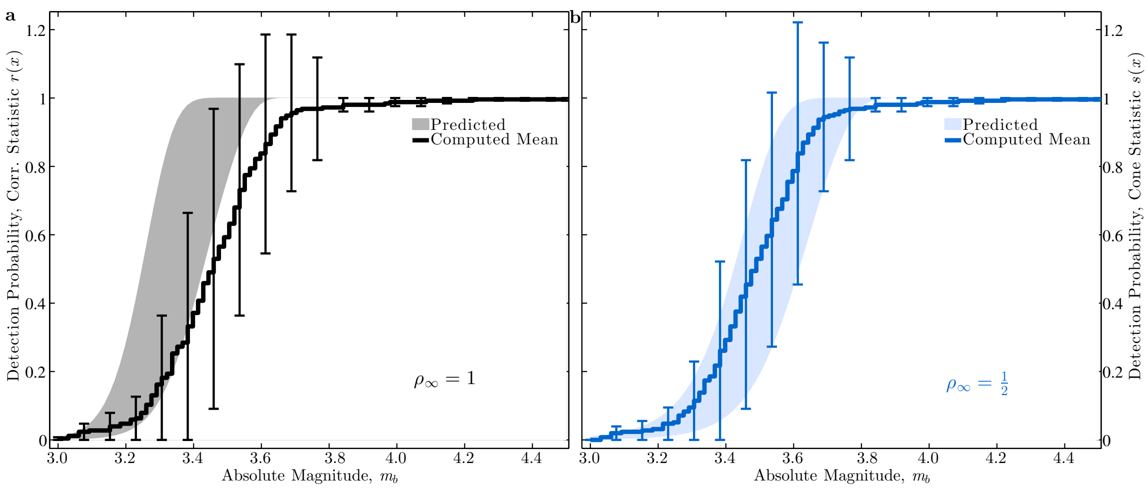

Figure 9: Observed and predicted receiver operating characteristic curves for $r(\boldsymbol{x})$ and $s(\boldsymbol{x})$ versus semi-empirical explosion magnitude. **a**: Shaded region shows range of ROC curves for $r(\boldsymbol{x})$ (Equation 3) that give the predicted detection performance in noise conditions recorded over 24 hrs on 08-Oct-2006. Superimposed stair plot shows the empirical detection performance (recorded detections/total events) averaged over 24 hr of data like that included in Figure 5. Error bars indicate the range in observed detection probability over the day and should not be misinterpreted as indicating that probability values exceeding one. **b**: Shaded region shows range of ROC curves for $s(\boldsymbol{x})$ (Equation 16) that give the predicted detection performance for the cone detector. Superimposed stair plot show observed detection performance averaged over 24 hr of data analogous to that shown in **a**.

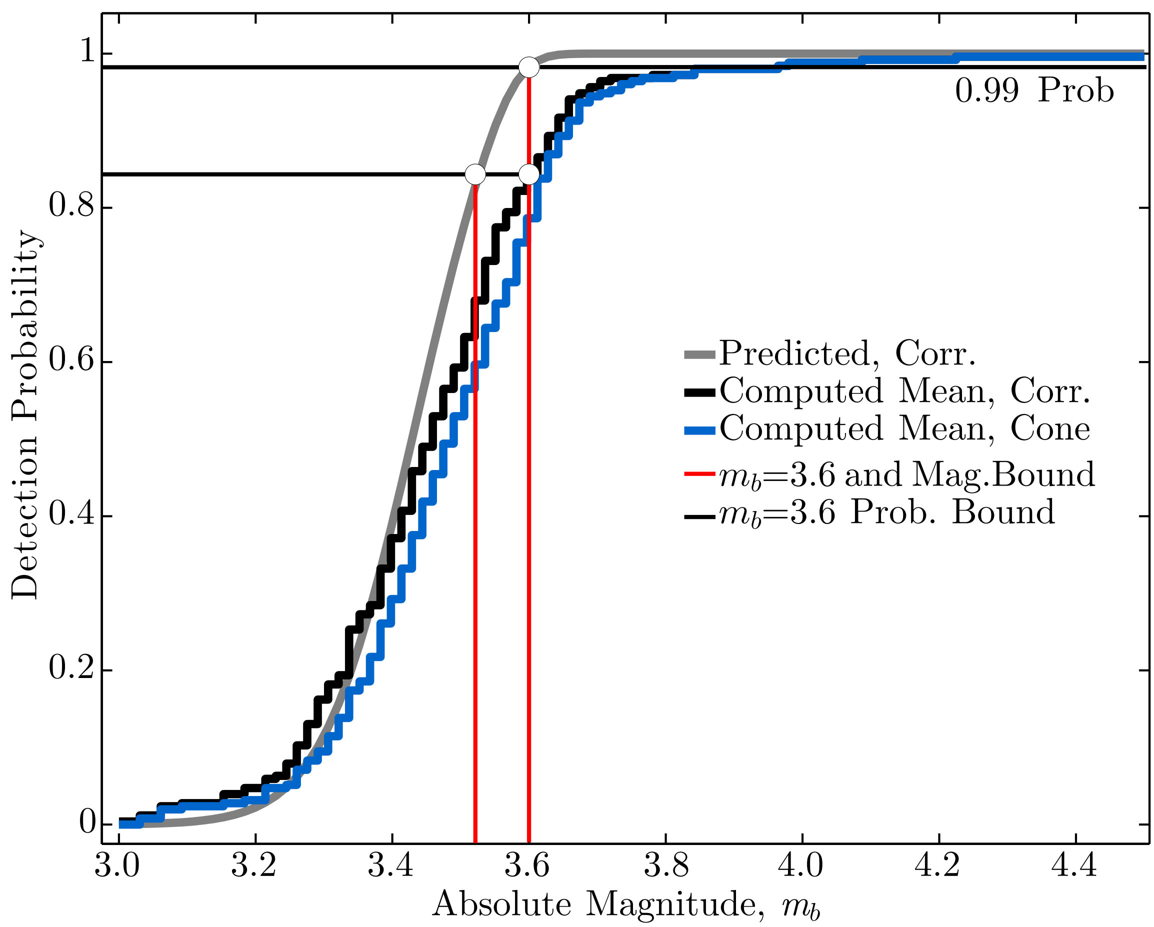

Figure 10: Superimposed empirical correlation and cone detection ROC curves from Figure 9. The correlation detector provides a marginally higher performance compared with the cone detector. The uppermost solid curve compares the lower-bound predicted correlation detector performance. The topmost horizontal line shows a reference 0.99 probability that intersects the prediction curve at the topmost circular marker. The corresponding vertical line shows the 0.99 probability explosion magnitude and its intersection with the observed correlation detection performance. The lowermost horizontal line shows the corresponding magnitude discrepancy at the predicted 0.99 detection probability (intersection marked by leftmost circular marker). The range between the horizontal lines that intersect the vertical axis measures the detection probability discrepancy.

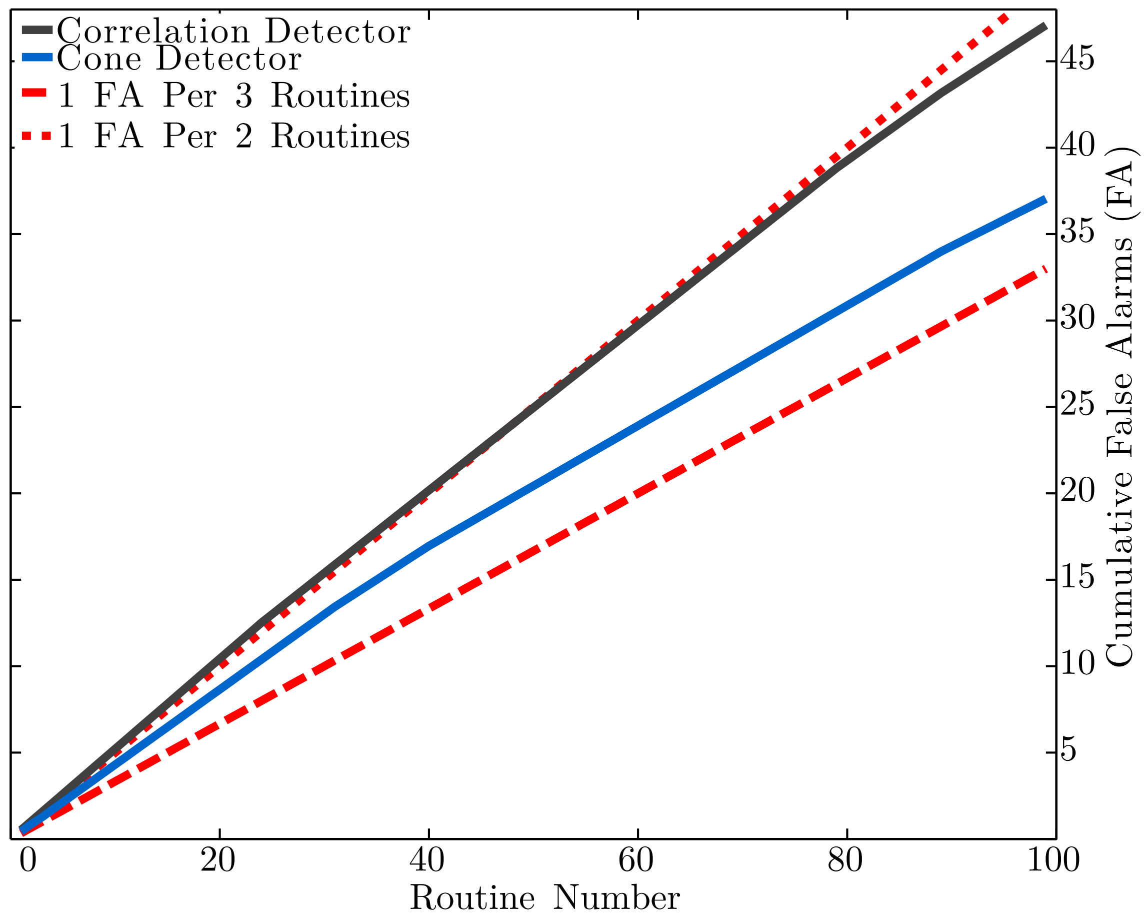

Figure 11: A comparison between cumulative false detection counts, per detector processing routines. The lowermost dashed curve shows a constant rate of one false alarm per three processing routine. The lowermost solid curve shows the observed number of false cone detections. The uppermost solid curve shows the observed number of false correlation detections. The uppermost dashed curve shows a constant rate of one false alarm per two routines.

Appendix A

This appendix develops the PDF and detection performance of the sample correlation detection statistic $r\left(\boldsymbol{x}\right)$ (Equation 3). Because the form of the sample correlation PDF differs from that reported elsewhere, we derive the general form here.

The relative square error in approximating a data stream $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ (the $\mathcal{H}_1$ model of Equation 15), with a waveform template $\boldsymbol{w}$, is the ratio of the least-squares error $||\boldsymbol{e}||^2$ to measure signal energy $||\boldsymbol{x}||^2$. This error may be re-written as a ratio of quadratic forms:

$$\begin{aligned}
\frac{||\boldsymbol{e}||^2}{||\boldsymbol{x}||^2} &= \frac{||\boldsymbol{x} - \hat{A}\boldsymbol{w}||^2}{||\boldsymbol{x}||^2} \\
&= \frac{\|\boldsymbol{x} - \dfrac{\langle \boldsymbol{x},\, \boldsymbol{w}\rangle}{\|\boldsymbol{w}\|^2}\boldsymbol{w}\|^2}{||\boldsymbol{x}||^2} \\
&= 1 - \frac{\langle \boldsymbol{x},\, \boldsymbol{w}\rangle^2}{||\boldsymbol{w}||^2||\boldsymbol{x}||^2} \\
&= 1 - r^2\left(\boldsymbol{x}\right)
\end{aligned} \tag{A.1}$$

where $\hat{A}$ is the maximum likelihood estimate for template waveform amplitude. We now rewrite the right hand side of the second equality in Equation A.1 as a ratio of subspace projections:

$$\frac{\|\boldsymbol{x} - \dfrac{\langle \boldsymbol{x},\, \boldsymbol{w}\rangle}{\|\boldsymbol{w}\|^2}\boldsymbol{w}\|^2}{||\boldsymbol{x}||^2} = \frac{\|P_W^{\perp}\left(\boldsymbol{x}\right)\|^2}{\|P_W^{\perp}\left(\boldsymbol{x}\right)\|^2 + ||P_W\left(\boldsymbol{x}\right)||^2} \tag{A.2}$$

where W is the subspace $\mathrm{span}\left(\boldsymbol{w}\right)$, $W^{\perp}$ is the orthogonal complement to W, P_W is the projector onto W and $P_W^{\perp}$ is the projector onto $W^{\perp}$. The denominator follows from the Pythagorean identity for Hilbert Spaces. We define two noncentrality parameters from these terms:

$$\begin{aligned}\lambda &= \frac{\|P_W\left(\mathbb{E}\left\{\boldsymbol{x}\right\}\right)\|^2}{\sigma^2} = \frac{||P_W\left(\boldsymbol{u}\right)||^2}{\sigma^2} = \rho_\infty^2 \frac{\|\boldsymbol{u}\|^2}{\sigma^2} \\ \lambda^\perp &= \frac{\|P_W^\perp\left(\mathbb{E}\left\{\boldsymbol{x}\right\}\right)\|^2}{\sigma^2} = \frac{||P_W^\perp\left(\boldsymbol{u}\right)||^2}{\sigma^2} = \left(1-\rho_\infty^2\right) \frac{\|\boldsymbol{u}\|^2}{\sigma^2}\end{aligned} \tag{A.3}$$

where the expected value and linear-projection operators commute. We combine the previous three equations to rewrite $r^2\left(\boldsymbol{x}\right)$:

$$\begin{aligned}1-\left(1-r^2\left(\boldsymbol{x}\right)\right) &= \frac{\|P_W^\perp\left(\boldsymbol{x}\right)\|^2 + ||P_W\left(\boldsymbol{x}\right)||^2}{\|P_W^\perp\left(\boldsymbol{x}\right)\|^2 + ||P_W\left(\boldsymbol{x}\right)||^2} \\ &\quad - \frac{\|P_W^\perp\left(\boldsymbol{x}\right)\|^2}{\|P_W^\perp\left(\boldsymbol{x}\right)\|^2 + ||P_W\left(\boldsymbol{x}\right)||^2} \\ &= \frac{\|P_W\left(\boldsymbol{x}\right)\|^2}{\|P_W^\perp\left(\boldsymbol{x}\right)\|^2 + ||P_W\left(\boldsymbol{x}\right)||^2} \\ &\overset{d}{=} \frac{\chi_1^2(\lambda)}{\chi_1^2(\lambda) + \chi_{N_E-1}^2(\lambda^\perp)}\end{aligned} \tag{A.4}$$

where $\overset{d}{=}$ is distributional equality, $\chi_{N_E-1}^2(\lambda^\perp)$ is the noncentral Chi-square distribution with $N_E - 1$ degrees of freedom and noncentrality parameter $\lambda^\perp$, and $\chi_1^2(\lambda)$ is the noncentral Chi-square distribution with one degree of freedom and noncentrality parameter λ. From the definition of the Beta distribution Kay (1998):

$$r^2\left(\boldsymbol{x}\right) \sim \mathrm{B}\left(t, \frac{1}{2}, \frac{1}{2}(N_E-1); \lambda, \lambda^\perp\right) \tag{A.5}$$

where $\mathrm{B}\left(t, N_1, N_2, \alpha, \beta\right)$ is the doubly noncentral Beta distribution function. It is evaluated at t (with the same domain as r^2), has N_1 and N_2 degrees of freedom, and noncentrality parameters α and β. The presence or absence of a target signal is indexed by the hypothesis $\mathcal{H}_k$ on the data. Hypothesis $\mathcal{H}_0$ symbolizes that the data consist of only noise, whereas $\mathcal{H}_1$ signifies that the data consists of a signal plus noise. The scalar N_E denotes the effective number of independent samples within $\boldsymbol{x}$. When the data stream contains only noise, the hypothesis $\mathcal{H}_0$ is satisfied and r^2 has a central Beta distribution, where $\lambda^{\perp} = \lambda = 0$. In the presence of signal, the data stream $\boldsymbol{x}$ generally has non-zero projections $P_W(\boldsymbol{x})$ and $P_W^{\perp}(\boldsymbol{x})$ that are respectively onto and orthogonal to the noise-contaminated template data vector $\boldsymbol{w}$. In this case λ, $\lambda^{\perp} \neq 0$, and r^2 has doubly noncentral Beta distribution. If the target signal is an amplitude scaled copy of the template waveform, then $\boldsymbol{x} = A\boldsymbol{w} + \boldsymbol{n}$, $\lambda^{\perp} = 0$, and r^2 has a noncentral Beta distribution. The singly noncentral Beta distribution therefore provides an absolute upper bound on the detection performance of a correlation detector consistent with assumptions underlying $\mathcal{H}_1$.

We derive PDF for r from the density of r^2 using a variable transformation; we additionally consider values $r < 0$:

$$
\begin{aligned}
f_R\left(r\left(\boldsymbol{x}\right);\mathcal{H}_k\right) = \mathrm{B}&\left(r^2\left(\boldsymbol{x}\right);\,\frac{1}{2},\,\frac{1}{2}\left(N_E-1\right),\,\lambda,\,\lambda^{\perp}\right) + \\
\mathrm{B}&\left(-r^2\left(\boldsymbol{x}\right);\,\frac{1}{2},\,\frac{1}{2}\left(N_E-1\right),\,\lambda,\,\lambda^{\perp}\right).
\end{aligned} \tag{A.6}
$$

The form of this distribution function differs from that derived in similar applications by Weichecki-Vergara and others Weichecki-Vergara et al. (2001). In that case, the signal-present data stream was assumed to correlate sample-by-sample with the template waveform, and the test statistic had a Pearson-moment product distribution.

Appendix B

This appendix outlines a method to estimate the effective degrees of freedom N_E of a data stream $\boldsymbol{x}$ as $\hat{N}_E$. This estimate provides an ostensible alternative to a full covariance matrix $\boldsymbol{\Sigma} \neq \sigma^2\boldsymbol{I}$ for $\boldsymbol{x}$ that is generally required for temporally or spatially correlated data. Density functions for detection statistics that process $\boldsymbol{x}$ are then easily parametrized by the effective number of independent data stream samples N_E. This scalar theoretically equates to twice the time (T)

bandwidth (B) product ($2TB$) of the data stream over the temporal duration of the correlation template. Real data often show that $N_E \ll 2TB$, however. This occurs both naturally and through processing operations like bandpass filtering, which replace each sample with itself plus a weighted sum of its neighbors and thereby introduce intra-sample statistical dependence. To quantify the influence such correlation exerts on the shape of our detector's PDF curves, we implement an empirical estimator for N_E, denoted $\hat{N}_E$, to continuously update such PDF parameterizations (e.g., $f_R\left(r\,;\,\mathcal{H}_0\right)$. This estimator computes the sample correlation between the multichannel template waveform $\boldsymbol{w}$ and several hundred psuedo-random, commensurate data vectors drawn from non-intersecting segments of pre-processed, signal sparse data within $\boldsymbol{x}$ (see ((Weichecki-Vergara et al., 2001, Section 2.4))). We compute the sample variance $\hat{\sigma}_R^2$ of the resultant correlation time series using 99.9% of the data by excluding 0.01% of the extreme left and right tails of its histogram. This provides the needed statistic to estimate N_E:

$$\hat{N}_E = 1 + \frac{1}{\hat{\sigma}_R^2}. \tag{B.1}$$

We use $\hat{N}_E$ to parametrize $f_R\left(r\,;\,\mathcal{H}_k\right)$ (Equation A.6), compute detector thresholds $\hat{\eta}$, and quantify detector performance.

Appendix C

This appendix develops the cone detection statistic (Equation 17), its PDF, and illustrates its performance; it was originally introduced in the context of icequake detection Carmichael (2013).

We first reference the two competing hypotheses of Equation 15:

$$\begin{aligned}
\mathcal{H}_0: \quad & \boldsymbol{x} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 \boldsymbol{I}\right) \\
& \text{(noise present, } \sigma \text{ unknown)} \\
\mathcal{H}_1: \quad & \boldsymbol{x} \sim \mathcal{N}\left(\boldsymbol{u}, \sigma^2 \boldsymbol{I}\right) \\
& \text{(noisy target present, } \boldsymbol{u} \in \mathcal{C},\ \sigma \text{ unknown)}
\end{aligned} \tag{C.1}$$

We further assume that pre-processing operations (like filtering) and naturally occuring temporal data correlation reduces the number of statistically independent samples in $\boldsymbol{x}$ to $N_E < N$. Given this parameterization, the PDF under $\mathcal{H}_0$ is denoted as $f_0\left(\boldsymbol{x};\, \mathcal{H}_0\right)$ and the PDF under $\mathcal{H}_1$ as $f_1\left(\boldsymbol{x};\, \mathcal{H}_1\right)$, where:

$$\begin{aligned}
f_0\left(\boldsymbol{x};\, \mathcal{H}_0\right) &= \frac{1}{\left(2\pi\sigma^2\right)^{\frac{1}{2}\cdot N_E}} \exp\left[-\frac{||\boldsymbol{x}||^2}{2\sigma^2}\right] \\
f_1\left(\boldsymbol{x};\, \mathcal{H}_1\right) &= \frac{1}{\left(2\pi\sigma^2\right)^{\frac{1}{2}\cdot N_E}} \exp\left[-\frac{||\boldsymbol{x}-\boldsymbol{u}||^2}{2\sigma^2}\right], \quad \boldsymbol{u} \in \mathcal{C}
\end{aligned} \tag{C.2}$$

We estimate the unknown parameters under each model listed in Equation 15 and select the applicable hypothesis using a log-generalized likelihood ratio test (log-GLRT). This test evaluates the PDFs for the competing hypotheses at the maximum likelihood estimates of their unknown parameter values and then compares their log-ratio to a threshold value η. The test's decision rule is a conditional, scalar inequality:

$$\Lambda\left(\boldsymbol{x}\right) = \ln\left[\frac{\max\limits_{\sigma,\boldsymbol{u}}\{\, f_1\left(\boldsymbol{x};\, \mathcal{H}_1\right)\,\}}{\max\limits_{\sigma}\{\, f_0\left(\boldsymbol{x};\, \mathcal{H}_0\right)\,\}}\right] \quad \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \quad \eta, \tag{C.3}$$

The maximum likelihood estimate of the variance under each hypothesis is Scharf and Friedlander (1994); Kay (1993, 1998):

$$\begin{aligned}
\hat{\sigma}_1^2 &= \underset{\sigma}{\operatorname{argmax}}\{\, \ln\left[f_1\left(\boldsymbol{x};\, \mathcal{H}_1\right)\right]\,\} = \frac{||\boldsymbol{x}-\boldsymbol{u}||^2}{N} \\
\hat{\sigma}_0^2 &= \underset{\sigma}{\operatorname{argmax}}\{\, \ln\left[f_0\left(\boldsymbol{x};\, \mathcal{H}_0\right)\right]\,\} = \frac{||\boldsymbol{x}||^2}{N},
\end{aligned} \tag{C.4}$$

where the subscripts on each sample variance estimate in Equation C.4 indicate the applicable hypothesis and N is the number of samples in $\boldsymbol{x}$, not to be confused with N_E. To estimate the distribution mean under $\mathcal{H}_1$, we evaluate $f_1(\boldsymbol{x}; \mathcal{H}_1)$ at the MLE for σ_1, and perform a constrained maximization of $\boldsymbol{u} \in \mathcal{C}$:

$$\begin{aligned} \hat{\boldsymbol{u}} &= \underset{\boldsymbol{u} \in \mathcal{C}}{\operatorname{argmax}} \{ \ln [f_1(\boldsymbol{x};\, \mathcal{H}_1)] \}|_{\sigma = \hat{\sigma}_1} \\ &= \underset{\boldsymbol{u} \in \mathcal{C}}{\operatorname{argmin}} \left\{ ||\boldsymbol{x} - \boldsymbol{u}||^2 \right\}. \end{aligned} \tag{C.5}$$

The solution to this equation,

$$\hat{\boldsymbol{u}} = P_{\mathcal{C}}(\boldsymbol{x}), \tag{C.6}$$

is the MLE for $\boldsymbol{u}$. That is, $\hat{\boldsymbol{u}}$ is the vector that minimizes the distance between the observed data $\boldsymbol{x}$ and all points that constitute the multiplet set $\mathcal{C}$. This defines the projection of $\boldsymbol{x}$ onto $\mathcal{C}$ as the unique signal that is either interior to, or on the boundary ∂C of $\mathcal{C}$ Stark and Yang (1998). The sample variance and cone-element MLEs from Equation C.4 reduce $\Lambda(\boldsymbol{x})$ to:

$$\begin{aligned} \frac{2}{N_E}\Lambda(\boldsymbol{x}) &= \ln\left(\hat{\sigma}_0^2\right) - \ln\left(\hat{\sigma}_1^2\right) \\ &= -\ln\left[\frac{||\boldsymbol{x} - P_{\mathcal{C}}(\boldsymbol{x})||^2}{||\boldsymbol{x}||^2}\right]. \end{aligned} \tag{C.7}$$

We obtained Equation C.7 without specifying $\mathcal{C}$ aside from it's convexity. This result therefore applies to any normally distributed data confined in a convex set. In the case that $\mathcal{C}$ is described by the correlation constraint of Equation 14, the projection energy has a conic decomposition:

$$||\boldsymbol{x} - P_{\mathcal{C}}(\boldsymbol{x})||^2 = ||\boldsymbol{x}||^2 - ||P_{\mathcal{C}}(\boldsymbol{x})||^2, \tag{C.8}$$

as given by Moreau's Decomposition Theorem Moreau (1962), which is analogous to the orthogonal subspace decomposition from linear analysis Stark and Yang (1998). The log-ratio is then expressible as:

$$\frac{2}{N_E}\Lambda\left(\boldsymbol{x}\right) = -\ln\left[1 - \frac{\left|\left|P_{\mathcal{C}}\left(\boldsymbol{x}\right)\right|\right|^2}{\left|\left|\boldsymbol{x}\right|\right|^2}\right]. \tag{C.9}$$

The right hand side of Equation C.7 is of the form $-\ln\left(1 - x^2\right)$, which is monotonically increasing for $0 \leq x \leq 1$. Since $\left|\left|P_{\mathcal{C}}\left(\boldsymbol{x}\right)\right|\right|^2 / \left|\left|\boldsymbol{x}\right|\right|^2 \leq 1$, Equation C.9 may be inverted for it's argument to provide an equivalent test statistic $s\left(\boldsymbol{x}\right)$ for the decision rule introduced in Equation C.3:

$$\frac{\left|\left|P_{\mathcal{C}}\left(\boldsymbol{x}\right)\right|\right|}{\left|\left|\boldsymbol{x}\right|\right|} = s\left(\boldsymbol{x}\right) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta. \tag{C.10}$$

Equation C.10 demonstrates that $s\left(\boldsymbol{x}\right)$ compares the ratio of the projected signal energy to the original signal energy. The exact projection $P_{\mathcal{C}}\left(\boldsymbol{x}\right)$ of $\boldsymbol{x}$, from a general Hilbert Space and onto $\mathcal{C}$, is derived in Stark and Yang (Stark and Yang, 1998, pages 111-113). Vector $P_{\mathcal{C}}\left(\boldsymbol{x}\right)$ is a nonlinear projection that depends on the value of $\boldsymbol{x}$, and is either in $\mathcal{C}$, on it's boundary $\partial\mathcal{C}$, or zero. We document an equivalent form of that projection here:

$$P_{\mathcal{C}}\left(\boldsymbol{x}\right) = \begin{cases} \boldsymbol{0}: & \dfrac{r}{\sqrt{1-r^2}} \leq -c, & P_{\mathcal{C}}\left(\boldsymbol{x}\right) = \boldsymbol{0} \\[2ex] \gamma\dfrac{\boldsymbol{z}}{\left|\left|\boldsymbol{z}\right|\right|}: & \dfrac{r}{\sqrt{1-r^2}} \in \left[-c, \dfrac{1}{c}\right], & P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C} \\[2ex] \boldsymbol{x}: & \dfrac{r}{\sqrt{1-r^2}} > \dfrac{1}{c}, & P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}. \end{cases} \tag{C.11}$$

The constants and vectors in Equation C.11 are:

$$\begin{aligned} r &\triangleq \frac{\langle\hat{\boldsymbol{w}}, \boldsymbol{x}\rangle}{\left|\left|\boldsymbol{x}\right|\right|} \\ c &\triangleq \sqrt{\frac{1-\rho_\infty^2}{\rho_\infty^2}} \\ \gamma &\triangleq \rho_\infty \left|\left|\boldsymbol{x}\right|\right| \left(r + c\sqrt{1-r^2}\right) \\ \frac{\boldsymbol{z}}{\left|\left|\boldsymbol{z}\right|\right|} &\triangleq \rho_\infty\hat{\boldsymbol{w}} + \sqrt{1-\rho_\infty^2} \cdot \frac{\boldsymbol{x} - \langle\hat{\boldsymbol{w}}, \boldsymbol{x}\rangle\hat{\boldsymbol{w}}}{\left|\left|\boldsymbol{x} - \langle\hat{\boldsymbol{w}}, \boldsymbol{x}\rangle\hat{\boldsymbol{w}}\right|\right|} \end{aligned} \tag{C.12}$$

The vector $\hat{\boldsymbol{w}}$ in Equation C.12 is the normalized multichannel template waveform defined by Equation 14, not to be confused with a parameter estimate. Using the definitions in Equation C.12, the test statistic $s(\boldsymbol{x})$ of Equation 16 is then expressible as:

$$s(\boldsymbol{x}) = \begin{cases} 0: & \frac{r}{\sqrt{1-r^2}} \leq -c, & P_{\mathcal{C}}(\boldsymbol{x}) = \mathbf{0} \\ \frac{\gamma}{||\boldsymbol{x}||}: & \frac{r}{\sqrt{1-r^2}} \in \left[-c, \frac{1}{c}\right], & P_{\mathcal{C}}(\boldsymbol{x}) \in \partial\mathcal{C} \\ 1: & \frac{r}{\sqrt{1-r^2}} > \frac{1}{c}, & P_{\mathcal{C}}(\boldsymbol{x}) \in \mathcal{C} \end{cases} \tag{C.13}$$

The detection statistic $s(\boldsymbol{x})$ must be equivalent to the correlation coefficient as ρ_∞ approaches a limiting value of one. To demonstrate this, we first note that any projected signal has a decreasing probability of lying inside the cone as ρ_∞ decreases (Equation C.13). Similarly, any projected signal has a decreasing probability zero of lying on the cone vertex. It follows that only the projection onto the limiting cone boundary $\partial\mathcal{C}$ is non-trivial, so that:

$$\begin{aligned} s(\boldsymbol{x}) &= \lim_{\rho_\infty \to 1^-} \frac{\gamma}{||\boldsymbol{x}||} \\ &= \lim_{\rho_\infty \to 1^-} \frac{\rho_\infty ||\boldsymbol{x}|| \left(r + c\sqrt{1-r^2}\right)}{||\boldsymbol{x}||} \\ &= r, \end{aligned} \tag{C.14}$$

where $\lim\limits_{\rho_\infty \to 1^-}$ is the limit of ρ_∞ approaching from values less than 1, whereby $c \to 0$.

We derive the requisite PDF for $s(\boldsymbol{x})$ from its CDF $F_S(s; \mathcal{H}_k)$ using the law of total probability. In words, this law states that the probability that the cone statistic random variable S takes on a value as large as $s(\boldsymbol{x})$ is the sum of three conditional probabilities: the probability that (1) $S < s(\boldsymbol{x})$, given $\Pr\{P_{\mathcal{C}}(\boldsymbol{x}) = \mathbf{0}\}$, plus (2) the probability $S < s(\boldsymbol{x})$, given $\Pr\{P_{\mathcal{C}}(\boldsymbol{x}) \in \partial\mathcal{C}\}$, plus

(3) the probability that $S < s(\boldsymbol{x})$, given $\Pr\{P_{\mathcal{C}}(\boldsymbol{x}) \in \mathcal{C}^{\circ}\}$:

$$\begin{aligned}
F_S\left(s\left(\boldsymbol{x}\right);\mathcal{H}_k\right) = \\
&\Pr\left\{S \leq s\left(\boldsymbol{x}\right) \mid P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}^{\circ}\right\} \cdot \Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}^{\circ}\right\} + \cdots \\
&\Pr\left\{S \leq s\left(\boldsymbol{x}\right) \mid P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right\} \cdot \Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right\} + \cdots \\
&\Pr\left\{S \leq s\left(\boldsymbol{x}\right) \mid P_{\mathcal{C}}\left(\boldsymbol{x}\right) = \mathbf{0}\right\} \cdot \Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) = \mathbf{0}\right\}
\end{aligned} \tag{C.15}$$

where $\Pr\{A|B\}$ is the conditional probability of A, given B is true and statement $P_{\mathcal{C}}(\boldsymbol{x}) \in \mathcal{C}^{\circ}$ is equivalent to $P_{\mathcal{C}}(\boldsymbol{x}) = \boldsymbol{x}$ ($\boldsymbol{x} \in \mathcal{C}$). To express $F_S(s;\mathcal{H}_k)$ in a computationally evaluable form, we write its three conditioning factors in terms of $r/\sqrt{1-r^2}$ through Equation C.13. Two of these terms are trivial to evaluate from the definition for $s(\boldsymbol{x}) = ||P_{\mathcal{C}}(\boldsymbol{x})|| \,/\, ||\boldsymbol{x}||$:

$$\begin{aligned}
\Pr\left\{S < s\left(\boldsymbol{x}\right) \mid P_{\mathcal{C}}\left(\boldsymbol{x}\right) = \mathbf{0}\right\} &= 0, \quad \text{since } s\left(\boldsymbol{x}\right) = 0 \\
\Pr\left\{S \leq s\left(\boldsymbol{x}\right) \mid P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}^{\circ}\right\} &= 1, \quad \text{since } s\left(\boldsymbol{x}\right) = 1
\end{aligned} \tag{C.16}$$

The other terms in Equation C.15 are expressible using $r/\sqrt{1-r^2}$ through Equation C.13:

$$\begin{aligned}
\Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right\} &= \Pr\left\{-c < \frac{r}{\sqrt{1-r^2}} < \frac{1}{c}\right\} \\
\Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}^{\circ}\right\} &= \Pr\left\{\frac{r}{\sqrt{1-r^2}} \geq \frac{1}{c}\right\}
\end{aligned} \tag{C.17}$$

We evaluate these probabilities by developing the CDF $F_Q(q;\mathcal{H}_k)$ for the ratio $q = r/\sqrt{1-r^2}$ (k = 0, 1). To do so, we make a change of variables with $F_R(r;\mathcal{H}_k)$:

$$F_Q\left(q;\mathcal{H}_k\right) = F_R\left(\frac{-\,q}{\sqrt{1+q^2}};\mathcal{H}_k\right) + F_R\left(\frac{q}{\sqrt{1+q^2}};\mathcal{H}_k\right) \tag{C.18}$$

where $F_R(r;\mathcal{H}_k)$ is the CDF for correlation statistic r. We then use Equation C.18 to write the identities from Equation C.17 in terms of $F_Q\left(q; \mathcal{H}_k\right)$:

$$\begin{aligned}
\Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right\} &= F_Q\left(\frac{1}{c}; \mathcal{H}_k\right) - F_Q\left(-c; \mathcal{H}_k\right) \\
\Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \mathcal{C}^{\circ}\right\} &= 1 - F_Q\left(\frac{1}{c}; \mathcal{H}_k\right)
\end{aligned} \tag{C.19}$$

Last, we evaluate the derivative of $\Pr\left\{P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right\}$ in Equation C.15 through a variable change on r, where $-c < r/\sqrt{1-r^2} < 1/c$. To do so, we first note that $s\left(\boldsymbol{x}\right) = \rho_\infty(r - c\sqrt{1-r^2})$ is invertible over this domain and express r as a function of s:

$$r(s) = \rho_\infty s - \sqrt{1-\rho_\infty^2}\sqrt{1-s^2} \tag{C.20}$$

The PDF of $s\left(\boldsymbol{x}\right)$ over $-c < r/\sqrt{1-r^2} < 1/c$ is therefore:

$$f_S\left(s; \mathcal{H}_k | P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}\right) = f_R\left(r\left(s\right); \mathcal{H}_k\right) \cdot \left|\frac{dr(s)}{ds}\right| \tag{C.21}$$

where $|P_{\mathcal{C}}\left(\boldsymbol{x}\right) \in \partial\mathcal{C}$ indicates the restricted domain of r. To obtain the PDF for $s\left(\boldsymbol{x}\right)$ over $-1 \leq r \leq 1$, we differentiate Equation C.15, and substitute Equation C.19 and Equation C.21 into the results. This produces a density function that depends on only s, c, and $f_R\left(\bullet; \mathcal{H}_k\right)$:

$$\begin{aligned}
&f_S\left(s; \mathcal{H}_k\right) = \\
&\left[F_Q\left(\frac{1}{c}; \mathcal{H}_k\right) - F_Q\left(-c; \mathcal{H}_k\right)\right] \cdot f_R\left(r\left(s\right); \mathcal{H}_k\right) \cdot \left|\frac{dr(s)}{ds}\right|
\end{aligned} \tag{C.22}$$

since the top line of Equation C.15 is constant and differentiates to zero.

We assessed the validity of our derivation using a Monte Carlo simulation whereby we projected random noise and noise-contaminated signal vectors onto several convex cones of increasing aperture and computed the statistic $s\left(\boldsymbol{x}\right)$. This simulation demonstrates that Equation C.22 agrees with our empirical histograms (Figure C.2).

We additionally compared our cone detector thresholds against constant false-alarm-on-noise constraints. We thereby inverted for cone detector thresholds using a fixed value for the effective degrees of freedom ($N_E = 400$) over a grid of decorrelation parameters $(1 - \rho_\infty)$ that ranged from 0 to 0.25. We repeated this process for for several false alarm rates (Figure C.3). The resulting detection thresholds increase most rapidly for small changes near $\rho_\infty = 1$, where the signal space of the cone geometrically collapses to the one dimensional subspace used by the associated correlation detector. The slope of the curves here may become undefined. We will explore this result more quantitatively in future work.

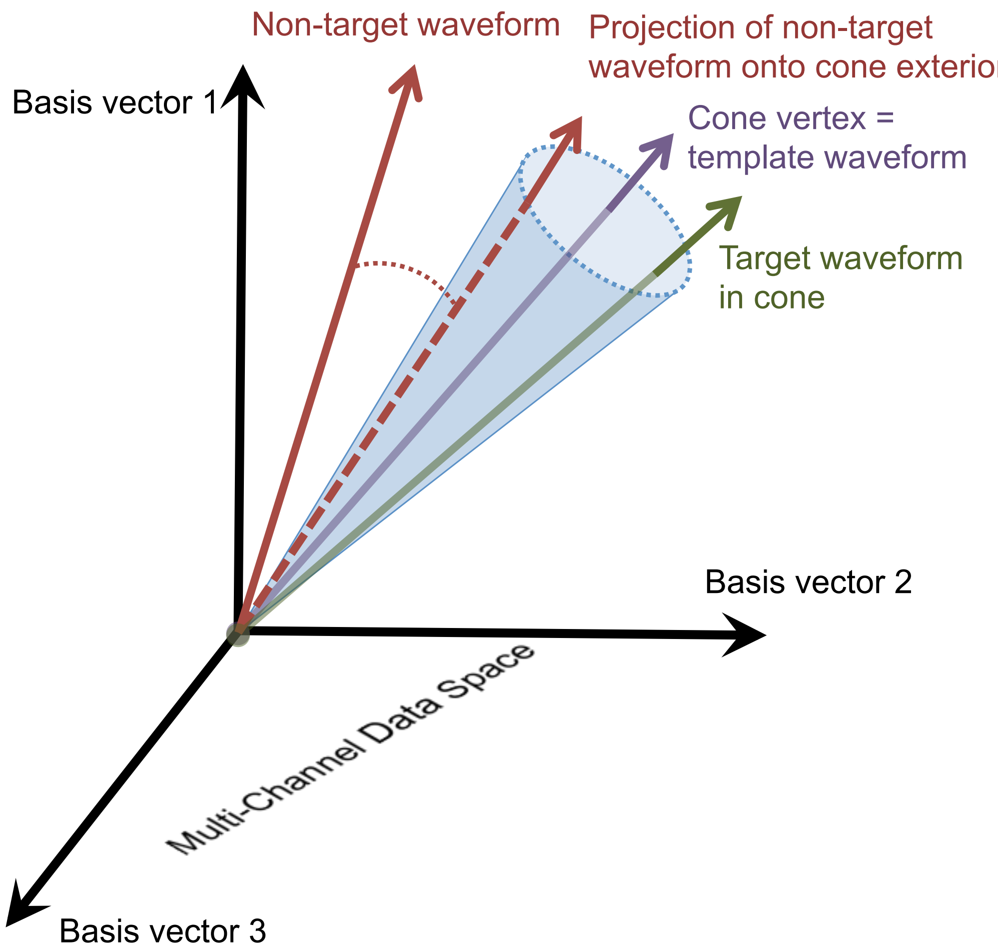

Figure C.1: Notional geometry of convex cone projections inside $\mathcal{C}$ and onto its boundary $\partial\mathcal{C}$. Multichannel signals are vectorially represented here in three dimensions but occupy hundreds or thousands of dimensions in practice.

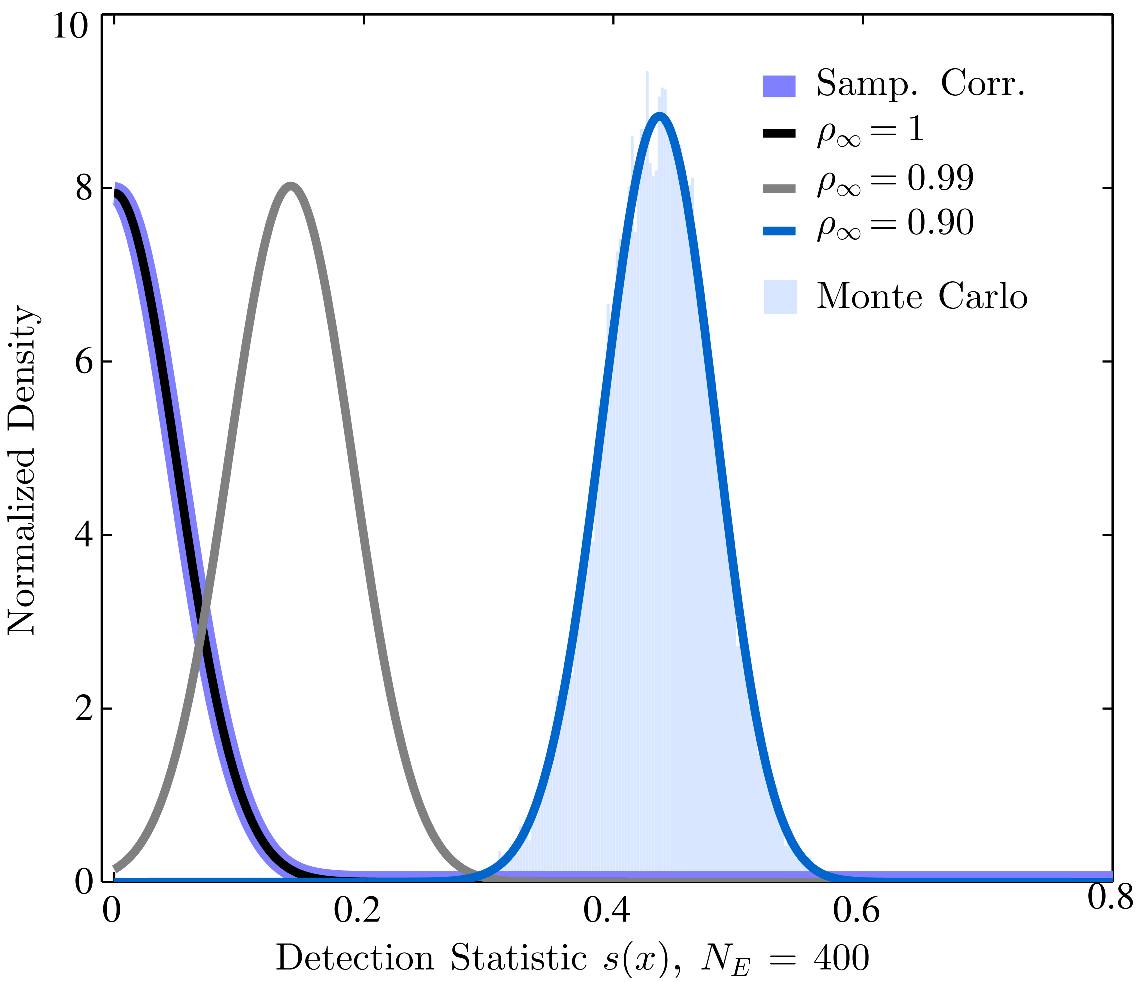

Figure C.2: Null hypothesis PDFs for three cases of deterministic template-target waveform correlation uncertainty: $\rho_\infty = 1$, $\rho_\infty = 0.99$, $\rho_\infty = 0.9$ (where $N_E = 400$). The PDFs for $r(\boldsymbol{x})$ and $s(\boldsymbol{x})$ equate for identically shaped waveforms as shown by the purple and black curves. The shaded region shows a Monte Carlo simulation of the PDF for $s(\boldsymbol{x})$ when $\rho_\infty = 0.9$.

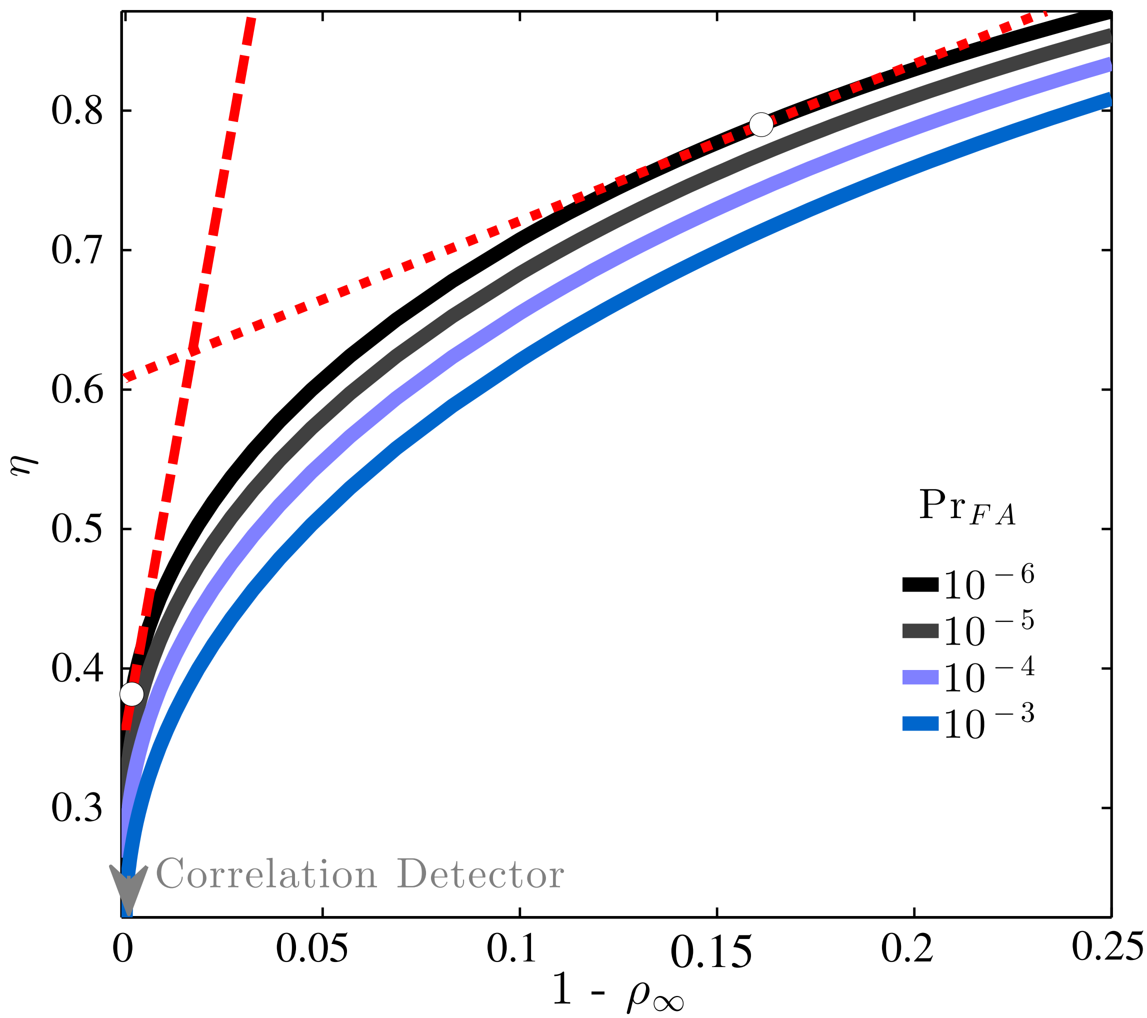

Figure C.3: Cone detector thresholds η (Equation 21) for constant values of Pr_{FA} compared against deterministic uncertainty in the template-target waveform cross correlation $1-\rho_\infty$, where $N_E = 400$. The left-most tangent line to the 10^{-6} curve shows a rapid increase in η for small uncertainties in deterministic template/target waveform correlation relative to near-linear increases in η for larger uncertainties. $\rho_\infty = 0$ corresponds to the correlation detector.

Appendix D

An unbiased estimator for the relative magnitude Δm between an explosion that generates a noisy waveform $\boldsymbol{x} = \boldsymbol{u} + \boldsymbol{n}$ and an explosion that generates a detector template $\boldsymbol{w}$ is (Carmichael and Hartse, 2016, Equation B.8):

$$\Delta \hat{m} = \frac{1}{2} \log_{10} \left(\frac{||\boldsymbol{x}||^2}{||\boldsymbol{w}||^2} - \frac{\hat{\sigma}_1^2 N}{||\boldsymbol{w}||^2} \right), \tag{D.1}$$

where σ_1^2 denotes an estimate for noise variance in $\boldsymbol{x}$ (subscript denotes hypothesis $\mathcal{H}_1$ is true) and N is the number of samples in $\boldsymbol{x}$, not to be confused with N_E. To relate $||\boldsymbol{x}||^2$ to the noncentrality parameters λ and $\lambda^{\perp}$ that shape the PDF $f_S\left(s; \mathcal{H}_0\right)$, we again use the standard Pythagorean identity for Hilbert Spaces:

$$||\boldsymbol{x}||^2 = \|P_W\left(\boldsymbol{x}\right)\|^2 + ||P_W^{\perp}\left(\boldsymbol{x}\right)||^2 \tag{D.2}$$

where the subspace projection terms are defined in Equation A.2. Next, we use the definitions of $\boldsymbol{u}$, λ, and $\lambda^{\perp}$ to rewrite $||\boldsymbol{x}||^2$ as:

$$||\boldsymbol{x}||^2 = \sigma_1^2 \left(\lambda + \lambda^{\perp}\right) + (\text{noise term}) \tag{D.3}$$

where "noise term" is an inner product expression that includes $\boldsymbol{n}$ and σ_1^2 is the true variance of the target data. The noncentrality parameters in Equation D.3 are related by (Carmichael and Hartse, 2016, Equation A.11):

$$\lambda^{\perp} = \left(\frac{1 - \rho_\infty^2}{\rho_\infty^2} \right) \lambda, \tag{D.4}$$

and the expected value of "noise term" in Equation D.3 is:

$$\begin{aligned}\mathbb{E}\left\{\text{noise term}\right\} &= \mathbb{E}\left\{2\langle \boldsymbol{u}, \boldsymbol{n}\rangle\right\} + \mathbb{E}\left\{||\boldsymbol{n}||^2\right\} \\ &= \sigma_1^2 N.\end{aligned} \tag{D.5}$$

where $\mathbb{E}\left\{\langle \boldsymbol{u}, \boldsymbol{n}\rangle\right\} = 0$ since the noise is zero mean. Finally, we combine the preceding equations (Appendix D), remove the noise-bias term $\sigma_1^2 N$, and write λ in terms of relative magnitude:

$$\lambda = \rho_\infty^2 \frac{||\boldsymbol{w}||^2}{\sigma_1^2} \cdot 10^{2\Delta m} \tag{D.6}$$

Equation D.6 thereby expresses the noncentrality parameter for both the correlation and cone detection statistic in terms of the four fundamental scalars describing the wavefield: the deterministic correlation ρ_∞ between the template and target waveforms, the signal energy $||\boldsymbol{w}||^2$ of the template waveform, the noise variance σ_1^2 of the target data, and the relative magnitude Δm between the template and target source.

For our purposes, Equation D.6 parameterizes λ by relative magnitude. In such cases, the term ρ_∞ is estimable from Equation 12, σ_1^2 is estimable from the top term of Equation C.4, and $||\boldsymbol{w}||^2$ is effectively deterministic since correlation detectors generally implement a high SNR templates.

Appendix E

We identified three potentially significant sources (risks) of error in our detection capability assessments and empirical comparison. The first risk of error is attributable to certain details of template selection. Specifically, waveforms recorded on IMS arrays with large differences in source-receiver separation distance show temporal gaps in start time of the high-SNR portions of the recorded signals and therefore require sample imputation. Zero padding such data to equalize length can lead to several biases, however. For example, supplementing data with a large number of zeros causes the empirical null PDF (histogram) for the correlation detection statistic to become more concentrated around its mean and thereby artificially lowers the detector's threshold. The statistic histogram then fits the predicted PDF poorly and biases estimates of the degree of freedom parameter $\hat{N}_E$, since the correlation variance is reduced by the presence of zeros (Equation B.1). To mitigate these problems and facilitate template scanning, we therefore abstained from zero-padding the intermediate, noisy portion of our waveform template. Instead, we shifted each seismogram channel by an amount equal to its start time, minus the earliest start time among all template channels (template and target data). We thereby maintained the signal-only length of our template, avoided padding, and kept all data aligned to the same origin time. In addition to preventing estimation bias, this shifting also mitigated needless computation of padded data and prevented divide-by-zero errors. We therefore consider our template selection to present a low (direct) risk of error.

Estimation of ρ_∞ presents a second potential source of error. This arises largely from ambiguous estimation schemes for waveform SNR that influence the variability of ρ_0 (Equation 11), which, in turn, inversely scales $\hat{\rho}_\infty$ (Equation 12). We mitigated ambiguity problems by carefully selecting an associated low variance estimate for SNR which normalizes ρ_∞ (Equation 12). One "common-sense" estimate for SNR is the ratio of an N-point sample variance estimate of the data proceeding the detected signal, divided by an N-point sample variance estimate of data preceding the detected signal (a renormalized STA/LTA). It is obvious, however, than any such estimate will be biased by background signals contaminating the data steam, which reduce the resultant SNR estimate for $\boldsymbol{u}$. It follows that $\hat{\rho}_\infty$ will be a biased estimator, and give lower-than-true deterministic correlation values. A better approach requires pre-processing target data with a power (STA/LTA) detector, removing samples that exceed a threshold for signal declaration, and then computing the noise variance $\hat{\sigma}_0^2$ from this remaining data. The quotient:

$$\hat{\mathrm{SNR}} = \frac{\|\boldsymbol{u}\|^2}{(N-1)\hat{\sigma}_0^2} \tag{E.1}$$

then gives a reduced-bias SNR estimate. We used this estimator to compute ρ_0 as:

$$\hat{\rho}_0 = \frac{\sqrt{\hat{\mathrm{SNR}}\left(\boldsymbol{w}\right)}}{\sqrt{\left(\hat{\mathrm{SNR}}\left(\boldsymbol{w}\right)+1\right)}}\frac{\sqrt{\hat{\mathrm{SNR}}\left(\boldsymbol{u}\right)}}{\sqrt{\left(\hat{\mathrm{SNR}}\left(\boldsymbol{u}\right)+1\right)}} \tag{E.2}$$

We therefore consider our estimates of ρ_0 to be robust, and unlikely to induce significant uncertainty.

Last, noise variance estimation presents an additional risk of error. As we noted before (see Estimation Deterministic Decorrelation), background seismicity adds signal to the time series and can thereby bias such estimates. We therefore processed our target data in parallel with a power (STA/LTA) detector that operated at a 10^{-6} false-alarm-on-noise probability, removed data that exceeded the associated threshold, and used the remaining (almost) signal free data to estimate noise variance on each IMS channel. This scheme thereby avoided bias from signal. A second form of bias was also present, however. This latter bias originated from the non-uniform number of channel samples processed by the detector. Specifically, we noted above that our template included large temporal gaps over portions of the waveform, and consequently, records from MJAR contribute less to each detection statistic than do longer duration records at MKAR. An alternative noise variance estimate that accounts for such disparate channel contribution employs the pooled variance $\hat{\sigma}_P^2$, given by:

$$\hat{\sigma}_P^2 = \frac{\sum_{k=1}^{L}(N_k - 1)\hat{\sigma}_k^2}{N - L} \tag{E.3}$$

where $\hat{\sigma}_k^2$ is the noise variance estimate from channel k. We performed a subsequent analysis using this estimator, and found that it was often smaller than our naive estimation that used equally weighted data. This would have increased the effective size of the noncentrality parameter λ that controls predicted detection power. It may explain why, at times, our observed detection capability exceeded the concurrent predicted detection capability. We therefore suggest that performance discrepancy may generally result from an inconsistency in noise variance estimation.